Behavioral Causes of the Bullwhip Effect: “Satisficing” Policies with Limited Information Cues

Rogelio Oliva
Mays Business School
Texas A&M University
Wehner 301F – 4217 TAMU
College Station, TX 77843
Tel (979) 862-3744; Fax (979) 845-5653
roliva@tamu.edu

Paulo Gonçalves
MIT Sloan School of Management
30 Wadsworth St. E53-339
Cambridge, MA 02142
Tel (617) 253-3886; Fax (617) 253-7579
paulog@mit.edu

Behavioral Causes of the Bullwhip Effect: “Satisficing” Policies with Limited Information Cues

We evaluate, in an experiment with the Beer Distribution Game, a complementary behavioral source of the bullwhip effect that has been previously ignored in the literature: overreaction to backlogs. By separating the estimation of the response to inventory and backlog, we find that players treat backlog differently than inventory. Contrary to our expectations, players do not over-order when in backlog; instead, they have a measured response, saturating order adjustment and limiting the amount of amplification they introduce in the order stream. We also find stronger evidence than previous studies that players underestimate the supply line, leading to a more unstable ordering policy. Using a simulated order stream, we find that players display bounded rationality and that their estimated decision policy is not different in form and performance than a policy that, with the information cues available in the Beer Distribution Game (inventory position and orders), minimizes local cost. Players’ estimated ordering policy, however, aims to maintain higher inventory levels, leading to increased order amplification and costs for upstream echelons. Hence, the estimated ordering policy presents strong behavioral components: it ignores the supply line and under-reacts to backlog while aiming for higher than necessary inventory levels. We conclude by discussing the implications of these findings for future research and practice in supply chain management.

Keywords: supply management, experiments, heuristics, system dynamics.

1. Introduction

The *bullwhip effect*, the tendency for the variability of orders to increase as one moves from customers to manufacturers, is a frequent and costly problem in supply chains, leading to excessive capital investment, inventory gluts, low capacity utilization and poor service (Armony and Plambeck, 2005; Gonçalves, 2003; Lee, Padmanabhan, and Seungjin, 1997a; Sterman, 2000). For instance, Hewlett-Packard lost millions of dollars in unnecessary capacity and excess inventory following a post-shortage demand surge for its LaserJet printers (Lee, Padmanabhan, and Seungjin, 1997b). Cisco Systems incurred more than US$ 2 billion inventory write-off due to a strong inventory built up followed by a drastic decrease in retailer orders (Adelman, 2001; Armony and Plambeck, 2005). The sources for this amplification in demand variability include operational causes — such as batching of orders, order gaming due to shortages, forward buying due to price discounts, and errors in demand forecasting (Lee *et al.*, 1997a) — and behavioral ones — such as failure to account adequately for the supply line of unfilled orders (Sterman, 1989) and the adoption of coordination stocks (Croson, Donohue, Katok, and Sterman, 2005).

Motivated by Sterman (1989; 1992), a number of experimental studies have used the Beer Distribution Game (BDG) to explore behavioral causes for the bullwhip effect and methods for dampening it (see Croson and Donohue, 2002, for a review.) Kaminsky and Simchi-Levi (1998) find that reducing the ordering and shipment lags decrease overall supply chain costs even though order amplification remains the same. Gupta, Steckel and Banerji (2001) and Steckel, Gupta and Banerji (2004) also show that reducing lead times leads to lower costs. Their results suggest, however, that the impact of sharing POS data on costs depends on the nature of customer demand. Using a stationary and known demand (as proposed by Chen and Samroengraja, 1999), Croson and Donohue (2003) find that POS data significantly reduces order oscillation — particularly in upstream echelons — and reduces overall supply chain costs. Croson and Donohue (2006) find a similar result when echelons shared inventory information.

Interestingly, the bullwhip effect occurs even when demand is *fixed*, commonly-known, and players start at the *optimal* inventory level (Croson *et al.*, 2005). The authors suggest that players build inventory to protect against coordination risk (i.e., the risk that others will deviate from optimal behavior.) Investigating the effect of learning and communication on the bullwhip effect, Wu and Katok (2006) find that order variability decreases when team players are allowed to formulate strategies collaboratively. All these studies, when estimating the ordering decision rule for individuals, arrive at a common source of supply chain instability: players underestimate the supply line of unfilled orders. This work contributes to this line of empirical research by articulating and analyzing a complementary behavioral source of the bullwhip effect that has been overlooked by previous research: overreaction in response to shortages. Overreaction implies that subjects order more aggressively (e.g., have a stronger reaction) when they face shortages than when they hold inventory.

As Mitchell suggests, when competing with other retailers for scarce supplies (i.e., horizontal competition), retailers inflate their orders to manufacturers to improve their chances of obtaining the supply they need.

> [R]etailers find that there is a shortage of merchandise at their sources of supply. Manufacturers inform them that it is with regret that they are able to fill their orders only to the extent of 80 percent. … Next season, if [retailers] want 90 units of an article, they order 100, so as to be sure, each, of getting the 90 in the pro rata share delivered." (1924, p. 645)

Since there is no horizontal competition in the BDG, we cannot justify overordering as a rational consequence of the rationing game proposed by Lee et al. (1997a). However, we hypothesize that players could overreact in response to backlogs motivated by Tversky and Kahneman's (1974) availability heuristic (i.e., the tendency to overreact to dramatic or vivid events). A backlog is a dramatic event in the beer game because (a) is twice more costly than inventory and (b) causes great disruption to the supply chain. Contrary to our expectations, we find that players do not overreact when in backlog, instead their correction saturates at a maximum value; a policy that is more stable than the linear response to backlogs suggested in previous studies. We also find stronger evidence than previous studies that players underestimate the supply line, leading to a more unstable ordering policy. Using a simulated order stream to test the rationale of these two components of the ordering policy, we found that players show bounded rationality using only the information available to them in a policy that is not significantly different in form and cost performance from the policy that minimizes local cost. The estimated ordering policy, however, aims to maintain a higher inventory level than the cost-minimizing rule, which leads to increased order amplification and costs for upstream echelons. Hence, the estimated ordering policy indicates a strong behavioral component to supply chain instability, i.e., it ignores the supply line and underreacts to backlog while aiming for higher than necessary inventory level.

The remainder of the paper is structured as follows. In §2 we present the experimental design and methods, in §3 our models and results. In §4 we perform sensitivity analysis of the parameters of the estimated rule and compare the estimated rule with the local cost minimizing rule with the available information cues. We conclude with a summary of our findings and implications for practitioners and researchers in supply chain management.

2. Experimental Design

Our experiment utilizes a web-based version of the Beer Distribution Game (BDG) developed at Harvard Business School that maintains the essential structure of the board game (Sterman, 1989). The game represents a serial supply chain with four echelons: retailer, wholesaler, distributor, and factory (R, W, D, and F, respectively). Each supply chain is independent of the other and managed by a team charged with minimizing the supply chain cost. Each echelon incurs an inventory holding cost of $0.50 per unit/week and a backlog cost of $1.00 per unit/week. Shipment and order delays between echelons are two weeks and factories incur a one-week production delay with no capacity constraints. Each simulated week players face the following sequence of events: (1) receive shipments; (2) fill customer orders, if sufficient inventory is available, otherwise accumulate a backlog; and (3) place an order with its supplier, where orders are constrained to be non-negative, i.e., it is not possible to cancel orders with the supplier.

The game is initialized in flow equilibrium: order and shipment flows are 4 units/week and each echelon starts with an initial inventory level of 12 units. Subjects are not informed about the shape of demand. A single time increase in retailer orders (a step input) is introduced in the second period (week), bringing orders to 8 units/week. To avoid end-of-horizon behavior the experiment is announced to run for a simulated year, but is, in fact, terminated after 36 weeks. The web-based version, by virtue of its automatic computation of order receipts, incoming orders, shipments, and inventory-backlog levels, can be run with less time pressure than the board version of the game on which a facilitator imposes the pace. The automatic recording of transactional data avoids reporting errors, although data entry (i.e., "typing") errors are still possible. Because of the cascading effects to other players, we did not attempt to correct typing errors.

Our data set consists of a sample of 116 pairs of first-year MBA students that played the game as part of the introductory course in operations management. This student sample has similar characteristics of previous studies using the BDG (Croson and Donohue, 2002; 2003; 2005; 2006; Sterman, 1989; Wu and Katok, 2006). Students had incentives to minimize team cost; the game was a graded assignment with team performance as the major component of the grade. The winning team also received a token award – similar to Sterman (1989). The students were on average 27 years old and had about two years of work experience in diverse areas. Prior to the game, players received a five-page document describing the structure of game and the sequence of events they will be facing. Less than two percent of the players expressed prior knowledge of the game. Due to the reduced number of "experts," they were not excluded from our sample. Players were randomly assigned, in pairs, to echelons (R, W, D, or F) and teams. Team members interacted via a computer screen and, in contrast to the board version of the game, lacked both visual access to the state of the supply line and knowledge of who other teammates were. We eliminated four games from our original sample. These games included one or more players showing anomalous ordering behavior (consistently not ordering when in backlog or placing high orders when holding large inventory) suggesting they had misunderstood the stock management task. Our analysis is based on the remaining 25 games.

2.1. Methods

We treated the BDG's non-negativity constraint on orders as censored data. That is, we assumed that an order for zero could represent situations in which a subject wished to cancel a previously placed order (a negative order) but was restricted by the rules of the game to a minimum order of zero. Accordingly, we estimated our model using a tobit model (Tobin, 1958). Also, to estimate a decision rule that reflects the full range of observations available, we structured the data from the games as a panel (cross-sectional time-series data set) with individual players the cross-sectional unit (i) and week of decision the time index (t). In contrast to previous studies that estimate the decision rules at the individual level, our panel data structure increases the efficiency of the estimates and the representativeness of the resulting rule as it allows us to make estimations across individuals and echelons. There being no reason to suspect that individual differences can be captured by changes in the constant term, and subjects being clearly a sample from a larger population, we assumed random effects across individuals (Greene, 1997). Estimations were performed using Stata's (2003) implementation of the random-effects cross-sectional time series tobit model and we tested the significance of the model's panel-level variance component by comparing the regression to the results of a pooled tobit regression.

3. Estimation of ordering policies

In a setting similar to the BDG with stationary and commonly-know demand distribution Chen (1999) demonstrated that a base-stock policy — where orders placed equal those received — minimizes total supply chain cost and avoids demand amplification. Since the demand in our experiment is both non-stationary and unknown to players (a step increase), it is not possible for players to calculate an optimal strategy prior to the game. While there is no reason to expect a base-stock policy to be optimal, it provides a starting point to test players' ordering policies. With the base-stock policy as a starting point, we incrementally increase the complexity of the ordering policies to include information cues and heuristics that players might have used. This incremental approach allows us to explore the marginal contribution of each of the information cues available to the players.

3.1 Base-stock and adaptive base-stock policies

A model that estimates the ability of a base-stock policy to explain the variability in orders is given by:

$$O_{it} = MAX(0, \beta_L L_{it-1} + u_i + \varepsilon_{it}) \quad \text{(I)}$$

where, to be consistent with the BDG, orders are constrained to be nonnegative; L_{it-1} represents orders received by the *i*th subject in the last period (*t*-1); u_i is the random disturbance characterizing the *i*th subject, and ε_{it} is an additive disturbance term. According to Chen's (1999) base-stock policy, orders placed (O_{it}) must equal those received in the previous period (L_{it-1}), thus we expect β_L to be equal to one.

Model I in Table 1 shows the estimated parameters for the base model together with the model's log-likelihood value, significance (χ^2), R^2, and root mean percent error. The model is highly significant (p<0.001), explains 54% of the variance in orders, and differences among players do not contribute to explain unexplained variance in orders ($\rho = \sigma_u^2 \big/ \left(\sigma_u^2 + \sigma_\varepsilon^2\right) = 0.0$). While a base-stock policy (using a lag forecast as a predictor for orders) provides a good fit for players' ordering policy, the fact that the β_L coefficient is slightly greater than one (p=0.03 for $H_0 : \beta_L = 1$) suggests that the rule does not fully capture all the adjustments being made by the players.

For cases when demand is non-stationary, Graves (1999) proposes an adaptive policy that, like the base-stock policy, replenishes the demand just realized, but adjusts orders to changes in demand. He suggests that the adjustment on orders should be based on the difference between successive demand forecasts projected over the supply lead-time:

$$O_t = L_{t-1} + LT\left(\hat{L}_t - \hat{L}_{t-1}\right) \tag{1}$$

where L_{t-1} is the demand just realized, $\hat{L}_t$ the current demand forecast, and LT the lead-time for the supply chain to replenish orders. Assuming a simple lag forecast $\left(\hat{L}_t = L_{t-1}\right)$, we define G as the difference between successive forecast $G_{t-1} = \left(\hat{L}_t - \hat{L}_{t-1}\right) = \left(L_{t-1} - L_{t-2}\right)$. Adding the non-negativity constraint and the random disturbances, and expanding the panel notation, yields the following model:

$$O_{it} = MAX\left(0, \beta_L L_{it-1} + \beta_{LT} G_{it-1} + u_i + \varepsilon_{it}\right) \tag{Ia}$$

Under this model, we expect β_L to be equal 1, i.e., full replacement of past orders, and β_{LT} to be close to the supply lead-time in the BDG, i.e., four weeks. Model Ia in Table 1 shows the results of the estimation of this model. The coefficient for the forecast adjustment (β_{LT}) is not significantly different from zero, thus the results are identical to the base-stock policy. We obtain similar results when we use, as suggested by Graves (1999), an exponential smoothing process to forecast demand in the adaptive base-stock policy. While the coefficient for the forecast adjustment is significant $(\beta_{LT} = 0.21, S.E. = 0.04)$, its value is considerably smaller than the expected four weeks of supply lead-time, and the additional regressor only explains an additional 1% of the variance in orders $\left(R^2 = 0.55\right)$. We, thus, failed to find evidence for an order adjustment based on demand forecast.

--- Insert Table 1 about here ---

3.2 Stock management policy

We revise the base-stock policy to incorporate inventory and supply line adjustments suggested by Sterman (1989). Due to the difficulty in finding the optimal ordering policy in the traditional BDG, Sterman proposes a simple, self-correcting ordering heuristic that uses information locally available to the decision maker and presumes no knowledge of the structure of the system. Specifically, managers are assumed to size orders to (1) replace expected losses from stock, (2) reduce the discrepancy between desired and actual stock, and (3) maintain an adequate supply line of unfilled orders. The decision rule is formalized as:

$$O_t = \hat{L}_t + \alpha_S (S^* - S_t) + \alpha_{SL} (SL^* - SL_t) \quad (2)$$

where $\hat{L}_t$ represents the expected loss from the stock, S_t and SL_t the inventory and supply line positions at time *t*, S^* and SL^* the desired levels for stock and supply line, and the parameters α_S and α_{SL} the fractional adjustment rate for inventory and supply line, respectively.

Sterman (1989) assumed adaptive expectations for the formation of the expected loss according to an exponential smoothing process:

$$\hat{L}_t = \theta L_{t-1} + (1-\theta)\hat{L}_{t-1} \quad (3)$$

and obtained, for each player, maximum likelihood estimates for the simultaneous equations 2 and 3, subject to the constraints $0 \le \theta \le 1$ and $\alpha_S, \alpha_{SL}, S^*, SL^* \ge 0$. The joint estimation of these equations, however, has the potential of shifting variance between the stock replenishment and forecasting equations, eqs. 2 and 3 respectively. Lower values of θ make the forecast series more stable and shift the residual variance to the replenishment decision, thus potentially biasing its parameter estimates (Oliva, 2003).

We assume a simple lag forecast $(\hat{L}_t = L_{t-1})$, an implied $\theta = 1$ in the exponential smoothing model in equation 3, and an intuitive and plausible model of expectation formation (Kleinmuntz, 1993). The simple lag forecast assumption has been used before in empirical research with the BDG (Steckel *et al.*, 2004); generates a reasonable forecast (0% Median Absolute Percent Error (MdAPE) and 23% Mean Absolute Percent Error (MAPE)) for a series with a coefficient of variation of 0.61; and has a Root Mean Square Error (RMSE) only 9.5% higher than the optimal exponential smoothing forecast. The change in forecast, the introduction of the non-negativity constraint on orders, and the expansions for panel data and additive disturbances yields the model:

$$O_{it} = MAX(0, L_{it-1} + \beta_0 + \beta_S S_{it} + \beta_{SL} SL_{it} + u_i + \varepsilon_{it}) \quad \text{(II)}$$

where $\beta_0 = \alpha_S S^* + \alpha_{SL} SL^*$, $\beta_S = -\alpha_S$ and $\beta_{SL} = -\alpha_{SL}$ in equation 2.

The results of the estimation of model II are presented in Table 1. The model is highly significant and explains 58% of the order variance, a 4% improvement over the base stock model. In this case 2% of the unexplained variance in orders is explained by the differences among individuals ($\rho = \sigma_u^2 \big/ \left(\sigma_u^2 + \sigma_\varepsilon^2\right)$) and this panel-level variance is marginally significant when compared to the pooled tobit model (likelihood-ration test of $\sigma_u^2 = 0$). The estimated fractional inventory adjustment ($\alpha_S = -\beta_S = 0.11$) is highly significant and has the expected sign. The estimated fractional adjustment to the supply line ($\alpha_{SL} = -\beta_{SL} = 0.01$), on the other hand, is not statistically significant, suggesting that players *ignore* the supply line. The dampened responsiveness of the estimated adjustment policies when compared to Sterman's (1989) previous estimations of the stock management model for individual players ($\overline{\alpha}_S = 0.26$ and $\overline{\alpha}_{SL} = 0.09$) suggests that our assumption of the simple lag-forecast is leaving less variance for the stock adjustment policies to explain.

3.3 Response to backlog policy

Previous estimations of ordering policies (Croson and Donohue, 2006; Sterman, 1989) treat backlog as negative inventory and assume a linear response to the gap between current and desired inventory. Because the cost of backlog is twice the holding cost for inventory, it is possible that subjects reacted differently to backlog than to excess inventory. To test, with the simplest model possible, for the possibility of a different reaction to backlog, we assumed a piecewise linear model (Pindyck and Rubinfeld, 1998), introducing a dummy variable (B_t) to reflect the backlog condition ($B_t = 1$ if $S_t < 0$; 0 otherwise). Accordingly, the response to inventory is modified to $\alpha_S\left(S^* - S_t\right) + \alpha_B S_t B_t$, where α_S represents the fractional adjustment rate for the inventory and α_B the incremental adjustment due to backlog (i.e., the response to backlog is $-\alpha_S + \alpha_B$). Since $S_t B_t \leq 0$, $\alpha_B < 0$ indicates a stronger reaction to the backlog condition (see Figure 1 for expected response to the inventory position and reaction to shortages). Retaining the assumption of a simple lag as the expected loss from the stock, the introduction of the non-negativity constraint on orders, and the expansions for panel data and additive disturbances lead to the following model:

$$O_{it} = MAX(0, L_{it-1} + \beta_0 + \beta_S S_{it} + \beta_B S_{it} B_{it} + u_i + \varepsilon_{it}) \qquad \text{(III)}$$

where $\beta_0 = \alpha_S S^*$, $\beta_S = -\alpha_S$, and $\beta_B = \alpha_B$.

--- Insert Figure 1 about here ---

The estimated model is highly significant (see model III in Table 1). The introduction of the backlog response explains another 2% of the order variance, increasing R^2 to 0.60, and the estimated fractional inventory adjustment ($\alpha_S = -\beta_S = 0.21$) is significantly higher (i.e., more aggressive) than the estimate found without the backlog adjustment (model II), suggesting that indeed players treat inventory differently than backlog. We were surprised, however, to find a positive coefficient for the backlog response ($\beta_B = 0.21$). The combined response to a situation when inventory is in backlog ($\beta_S + \beta_B$) is not significantly different from zero (see Backlog effect and its test in Table 1). This suggests that players in backlog place orders equal to the expected loss plus a constant amount (β_0) proportional the desired inventory level S^* (see Figure 2 for a schematic of the order response to the estimated model). Instead of "over-reacting" to a backlog situation as we had expected, players seem to ignore the backlog information cue and respond only to the $S_t = 0$ signal.

--- Insert Figure 2 about here ----

When the sample is split by echelon (see R, W, D, and F models in Table 2), model III is significant for all positions and all the estimates have consistent signs and are significant (the model with dummies for each position yields the same results as Model III and the coefficients for the three dummies are not significant). The seemingly paradoxical result that both the R^2 and RMSE of the models increase as we move up the supply chain is explained by the fact that each successive stage the supply chain faces a demand stream with higher variance (see $\bar{\sigma}_{L_{t-1}}$ row in Table 3) and, although the models explain a higher fraction of that variance (R^2), the magnitude of the errors (RMSE) is increasing. In terms of parameter estimates, as a result of the differences in variance in demand each echelon faces, the aggressiveness of the inventory fractional adjustment ($\alpha_S = -\beta_S$) increases as we move up the supply chain: the higher the variance of incoming orders (L), the more aggressive corrections to deviations in inventory. Parameter estimates for the three first echelons (R, W and D), however, are not statistically different and the pooled model for these non-factory echelons (~F in Table 2) yields more efficient estimates, all within the standard error of the estimates for the separate models. Factories, whose stock management problem is structurally different from that of the other echelons — their delivery delay is shorter, and, because they are uncapacitated, it remains constant — have a significantly different replenishment rule with fractional adjustments to inventory more than twice as aggressive as the other echelons. The decision rule is, however, consistent across echelons as the other parameters adjust to accommodate the required aggressiveness of the fractional inventory adjustment — estimated values for desired inventory ($S^* = \beta_0 / -\beta_S$) for each sample partition are not statistically different (see Table 3) and are consistent with one week of orders at the increased consumption rate — and maintain an almost-flat response to the backlog condition ($\beta_S + \beta_{SL} \approx 0$).

--- Insert Table 2 about here ---

--- Insert Table 3 about here ---

To test the robustness of the lack of response to backlog we tested each of the assumptions made on our model specifications. First, we tested the forecasting assumption and replaced the simple-lag forecast with static expectations ($\hat{L}_{it} = L^*$), extrapolative expectations ($\hat{L}_{it} = L_{it-1} + \gamma(L_{it-1} - L_{it-2})$, where $0 \le \gamma \le 1$), and optimal adaptive expectations ($\hat{L}_{it} = \theta_i L_{it-1} + (1-\theta_i)\hat{L}_{it-1}$, where θ_i minimizes the forecast error for each player). The almost-flat response to the backlog condition held under all forecasting assumptions. Second, we tested the non-linear shape of the response to inventory-backlog by testing different breakpoints for the piecewise linear model (the best fit was obtained when breakpoint is at *S*=0); introducing separate intercepts to the two line segments to decouple the inventory to the backlog response (the second intercept was not significant); and testing different non-linear continuous responses (e.g., quadratic and logistic models) — all models generated a flat response to the backlog condition and none was as intuitive as model III presented above. Third, we tested the censored data assumption and ran the model without the tobit constraint and found no significant change to the backlog response. Fourth, we tested the aggregation assumption and ran the model, first, as a pooled data set (as expected from the small value of ρ, there was no significant change to the estimates), and then for each player independently: 84% of the players showed underreaction to backlog, and the model was significant for 67% of those players. Finally, we tested the model with Sterman's (1989) data and an additional data sample and found similar results, ruling out explanations differences from our data collection methods. The next section explores the impact of the supply line on the almost-flat response to the backlog condition.

3.4 Response to supply line

Model IV (Tables 1 and 2) shows the estimated parameters and performance statistics for a decision rule that incorporates, in addition to the inventory and backlog adjustments, the supply line signal:

$$O_{it} = MAX(0, L_{it-1} + \beta_0 + \beta_S S_{it} + \beta_B S_{it} B_{it} + \beta_{SL} SL_{it} + u_i + \varepsilon_{it}) \qquad \text{(IV)}$$

For the full sample, the coefficient for the supply line is not significant and introducing it into the regression has no effect in the other estimates. When we split the sample in factories and non-factories, the β_{SL} coefficient becomes significant, taking the expected negative value for factories but a positive value for non-factories. Nevertheless, the introduction of the supply line as a regressor does not have a significant impact in the inventory and backlog adjustment fractions for the two sub-samples, and the models' overall performance does not improve. When estimating model IV for individual players only 21% of the non-factories and 52% of the factories showed a significant negative response to the supply line. These results indicate that players consistently *ignore* the supply line, a finding stronger than Sterman's (1989) that players underestimate the supply line.

A possible explanation for the positive and significant coefficient of the supply line for non-factories is that it reflects the net effect from a decrease in orders due to a large supply line (the expected negative coefficient) and an increase in orders due to a higher desired supply line (SL^*). We tested this by modeling an endogenous adjustment of the desired inventory and supply line levels. Assuming an endogenous goal formation as a function of the loss forecast ($S_t^* = \kappa_S + \gamma_S \hat{L}_t$ and $SL_t^* = \kappa_{SL} + \gamma_{SL} \hat{L}_t$) in Sterman's original equation (eq. 1); replacing $\beta_0 = \alpha_S \kappa_S + \alpha_{SL} \kappa_{SL}$, $\beta_S = -\alpha_S$, $\beta_{SL} = -\alpha_{SL}$ and $\beta_L = 1 + \alpha_S \gamma_S + \alpha_{SL} \gamma_{SL}$; assuming a simple lag forecast $(\hat{L}_t = L_{t-1})$; and reintroducing the backlog, panel and random variation terms, yields the linear model

$$O_{it} = MAX(0, \beta_0 + \beta_S S_{it} + \beta_B S_{it} B_{it} + \beta_{SL} SL_{it} + \beta_L L_{it-1} + u_i + \varepsilon_{it}).$$

Under the endogenous goal formation hypothesis $\alpha_S, \gamma_S, \alpha_{SL}$, and $\gamma_{SL} \geq 0$, thus we would expect $\beta_L \geq 1$. We rejected this explanation since this model for the non-factories in our sample yielded $\beta_L = 0.77$ (S.E. = 0.03), i.e., p=0.999 for H_0: $\beta_L < 1$.

4. Sensitivity Analysis

A decision rule that does not consider the supply line amplifies orders and increases overall chain cost. Because orders placed to correct for an inventory imbalance do not arrive instantaneously, on-hand inventory remains low and correction orders are placed again and again resulting, eventually, in an overshoot of the inventory level. However, a decision rule that has a flat response to backlog limits the amount of amplification it introduces to the order stream. To understand the effects of this limited response relative to the inherent amplification that results from ignoring the supply line, we explored the impact of different parameters of the players' decision heuristic on order amplification and cost.

4.1 Sensitivity to decision parameters

To test the sensitivity of the different elements of the stock replenishment decision rule we used a simulated order stream. We used correlated noise with $O_t = MAX(0, \sim N(\mu, \sigma, \tau))$ where $\mu = 8$ to reflect the steady state condition of the BDG; $\sigma = 4$ to approximate the incoming orders of our non-factory sample (see Table 3); and $\tau = 1.75$ to represent the estimated autocorrelation time constant from orders received by non-retailers in our sample. We used two indicators to evaluate the performance of the decision rule. First, to assess the replenishment decision rule's contribution to the bullwhip effect we measured *order amplification* as the ratio of the standard deviation of the outgoing order stream (the orders placed by the decision rule) to the standard deviation of the incoming order stream. Second, to capture the inventory and service levels that resulted from applying the decision rule we used the cost structure stipulated in the rules of the

BDG—\$0.50 unit/week for holding inventory and \$1.00 unit/week for backlog—to estimate the *average weekly cost*.

The simulated decision rule managed a position that faced, as in the BDG, a two-week information delay to communicate orders upstream and two-week transportation delay for orders to arrive. Because we assumed the supplier to have an infinite supply, any backlog incurred in this simulated environment is the result of structural delays and the decision rule. This assumption underestimates, relative to the BDG, the operating range for backlog for the simulated player. But because the decision rule does not take supply line into consideration and each decision made is based on the last period's order and current inventory/backlog position, the assumption does not affect the behavior of the decision rule per se and the results are comparable across simulations (tests with a model of the full supply line yielded qualitatively the same results).

Univariate sensitivity. To test the influence of each decision parameter we performed Monte-Carlo simulations, varying each parameter in isolation while maintaining the other parameters at the values estimated in the base case (Model III in Table 1; $\beta_0 = 1.69,\ \beta_S = -0.21,\ \beta_B = 0.21$). Each simulation was run for 300 weeks and since the model was initiated in equilibrium both measures of performance were calculated using data from the full simulation horizon. Figure 3 shows the resulting order amplification and average weekly cost of 1,000 simulations varying each parameter in the displayed range.

The β_0 parameter represents the constant order quantity players would place in the absence of inventory and, with β_S, determines the desired inventory level ($S^* = \beta_0 / -\beta_S$). Being a constant term this parameter has a slight impact on order amplification (Figure 3a). Higher values of β_0 increase order amplification somewhat as the decision rule now has “more room” to adjust to excess inventory: with orders constrained to be non-negative the maximum inventory downsizing adjustment is limited to $-L_{t-1}$ when $S_t \geq (\beta_0 + L_{t-1})/-\beta_S$ ($S_t \geq S^{'}$ in Figure 2); a larger β_0 increases $S^{'}$ and the operating range of the decision rule. As expected, the relationship between the target inventory level (S^*) and cost is u-shaped (Figure 3a')—lower values of desired inventory result in a higher probability of running into a backlog situation and higher levels of inventory in excessive carrying cost.

--- Insert Figure 3 about here ---

The fractional adjustment of the inventory position (α_S) is represented in the replenishment decision rule by $-\beta_S$. Note that the decision rule is not robust if the combined response to backlog ($\beta_S + \beta_B$) is positive — if $\beta_S + \beta_B > 0$ the order response to backlog becomes weaker as backlog grows, eventually shutting down orders and allowing backlog to grow, resulting in ever increasing weekly cost (see Figure 3) — thus, the relevant range for β_S is $(-\infty, -\beta_B]$. Panels b and b' in Figure 3 show order amplification and average weekly cost decreasing in β_S (i.e., as inventory correction becomes less aggressive) — more negative values of β_S represent lower target inventory ($S^* = \beta_0/-\beta_S$) increasing the probability of running into backlog.

Figure 3, panels c and c', captures the combined value of $\beta_S + \beta_B$ obtained by varying β_B while maintaining β_S at its estimated value. As in the case with inventory adjustment, amplification and cost are increasing in the intensity of the backlog adjustment. Although increasing the aggressiveness of the backlog adjustment increases the variability of the results, the average effect on performance is not as detrimental as when varying β_S. The minimum amplification and cost is obtained when $\beta_S + \beta_B = 0$, suggesting that for a decision rule that does not take supply line into consideration the best response to backlog is *no* response. Panels d and d' in Figure 3 show the sensitivity of the decision rule when varying both parameters simultaneously ($\beta_S = -\beta_B$). Variations in the response to inventory have little effect on amplification. The cost function, however, is highly sensitive to the effect that β_S has on the implied target inventory ($S^* = \beta_0 / -\beta_S$).

Multivariate sensitivity. To separate the mutual dependency among parameters in the decision rule we performed a sensitivity analysis, varying both β_0 and β_S while holding to the assumption of no backlog response ($\beta_B = -\beta_S$). Figure 4 shows the surfaces of average amplification and average weekly cost as functions of β_0 and β_S for 50 realizations of the incoming order stream. The general trends of the response surface are clear. Amplification is increasing in β_0 and in the intensity of the fractional adjustment of inventory, but there is a minimum cost basin at $\beta_0 \approx (0.5, 1.5)$ and $\beta_S \approx (-0.1, -0.2)$, values consistent to the ones estimated from our non-factory sample.

--- Insert Figure 4 about here ---

4.2 Comparison to local-cost-minimizing policy

To further assess the performance of the estimated policy we compared it to the local-cost-minimizing policy with the same information cues. For 50 different order realizations we found through a grid search of the parameter space the parameter values that would minimize cost. The average values for the cost-minimizing parameters β_0^*, β_S^*, and β_B^* were 0.89, -0.14, and 0.14, with standard errors of 0.06, 0.01, and 0.02, respectively. Confirming our interpretation of the sensitivity analysis, we found that for the local-cost-minimizing policy the response to backlog was not statistically different from 0 (p=0.339 for $H_0 : \beta_S^* + \beta_B^* = 0$). The response to the inventory position estimated from our non-factory sample ($-\beta_S = \beta_B = 0.17$) was slightly more aggressive than that of the local-cost-minimizing policy (p=0.003 for $H_0 : \beta_S = \beta_S^*$, and p=0.067 for $H_0 : \beta_B = \beta_B^*$), but the main difference between policies was that the estimated policy aimed for a higher inventory target ($S_e^* = 8.47 > S_o^* = 6.36$) to reduce the probability of backlog (p=0.001 for $H_0 : \beta_0 = \beta_0^*$).

Evaluating the performance of the policies to 50 different order realizations revealed the average cost for the estimated rule to be only 2.5% higher than the local-cost-minimizing policy, a non-significant difference (p=0.147 for $H_0 : c_o = c_e$). Order amplification, however, was found to be 4% higher for the estimated policy ($A_e = 1.12 > A_o = 1.07$), a highly significant difference (p=0.000 for $H_0 : A_o = A_e$). Note that the tests applied the estimated policy consistently throughout the simulation and that this consistency resulted in better performance than that of most players in our sample (Bowman, 1963). Nevertheless, these results suggest that players were quite accurate in devising a rule that would minimize local cost given the information cues available. This cost minimization, however, came at the expense of increasing variance in upstream orders.

5. Discussion

We explored, in an experimental serial supply chain, the causes of the bullwhip effect by proposing a complementary behavioral source of order amplification in supply chains: overreaction to backlogs. The paper contains several contributions relative to previous work in this area. Regarding our core hypothesis, we found that players treated backlog and inventory differently. Contrary to our expectations, however, players do not over-order when in backlog; instead, they have a measured response, saturating order adjustment at a maximum value and limiting the amount of amplification they introduce in the order stream. This result held across echelons and is robust across model specifications and data sets.

The fact that players do not account for the supply line clearly indicates that they are not fully rational optimizers. An ordering rule that fails to account for the supply line will lead to order amplification. Although the computation to estimate the supply line is simple (cumulative difference between orders placed and orders received), the form players use to track inventory status and orders placed in the BDG does not provide the space, nor is time given when playing the game, to perform this computation. The supply line information is not salient in the game cues, nor does it play a role in the cost function for the individual players; thus, it is not surprising that players tend to ignore it (Plous, 1993). Furthermore, players are not completely naïve considering that they do not over-order when in backlog. By not responding to the backlog condition, players create an ordering policy that is more stable than the linear response to inventory discrepancies. When evaluating the performance of the estimated rule, we found that players act as boundedly rational using the information available to them in a policy that is not significantly different in form and cost performance from the policy that minimizes local cost. The estimated rule, however, aims to maintain a higher inventory level than the cost-minimizing rule. Higher desired inventory increases order amplification as it creates more room for the decision policy to adjust inventory and increase the size of order adjustment. Higher order variance, in turn, increases costs for upstream echelons. Our findings suggest that boundedly rational players adopt a policy that "satisfices" local cost minimization at the expense of upstream order amplification and higher team costs. The estimated ordering policy indicates a strong behavioral component to supply chain instability, i.e., it ignores the supply line and underreacts to backlog while aiming for higher than necessary inventory level.

In addition, refinements in assumptions and estimating techniques shed light on important aspects of the estimated decision rules for BDG players. By structuring the data as a panel (cross-sectional time series), we use all the data available for estimating the replenishment decision rule, thereby increasing the efficiency of estimates and the representativeness of the resulting rule. Our estimated ordering policies compare well to the policies suggested by previous studies with the BDG that estimate parameters for individual players (Croson and Donohue, 2003; 2006; Sterman, 1989). While Sterman's (1989) study yielded average R^2 and RMSE of 0.71 and 2.86 respectively, it required four parameters per player (i.e., 176 parameters for 44 players) to achieve this result. Our most parsimonious model achieves an R^2 of 0.60 and an RMSE of 4.46 with only three parameters for 100 players, confirming that it is possible to make some inferences across individuals. The panel data structure also allowed us to perform analyses by echelon in the simulated supply chain. We found that the inventory adjustment fraction became more aggressive as players faced increasing order variance (e.g., upstream in the supply chain) but they did not change their desired inventory levels. Finally, by decoupling the estimation of the forecast rule from the stock replenishment rule, and removing constraints in the feasible space for all parameters, we found a weaker inventory correction response and even stronger evidence that players underestimate the supply line. Whereas previous research suggests that the supply line is under-accounted for, we only found a significant effect of the supply line in the decision rule for factories, which face shorter and uncapacitated supply lines.

Our findings have several implications for future research on ordering policies for supply chains. Our main finding, that decision-makers treat inventory and backlog differently, suggests a revision to behavioral models for replenishment decisions to account explicitly for the backlog condition. By assuming that people treat inventory and backlog similarly, models may overestimate their measures of managerial responses to inventory and backlog, potentially offering inappropriate managerial recommendations. Clearly, our finding of under-ordering when in backlog needs to be further explored since it indicates a larger departure from optimal performance. Among the potential explanations of this result are expectation adjustments that players could be making when faced with consistent under-delivery from suppliers – i.e., players might be adjusting the orders downward when faced when poor delivery performance from suppliers – and the small difference between inventory and backlog costs used in the BDG – i.e., the relative low penalty for running into backlog might not be enough to trigger and overreaction. We also found evidence that the current decision making limitation in the BDG seems to be access to and clear emphasis of relevant information cues. Previous research suggests that while access to relevant information cues is critical, access alone is insufficient to improve decision making in dynamic environments. Sterman (1987), Kampmann (1992), and Diehl and Sterman (1995) report that supply line information was readily available in their studies, but they were still severely underestimated. Richardson and Rohrbaugh (1990), however, find that clearly emphasizing important information cues, instead of just providing access to it, improves the outcome of the decision rules. The BDG, and most information systems supporting real supply chains, lack the proper mechanisms to make global indicators (e.g., the costs of amplifying upstream orders) more salient and more likely to be incorporated in the decision rule. Although it is now standard practice in supply chain design to account for supply line information and to increase inventory visibility throughout the chain, more salient local information cues and incentives shape the decision-makers policies to optimize locally, despite the detrimental impact that their decisions may have on the whole supply chain. The individual savings achieved by attempting to minimize local costs may, as in the BDG, greatly overshadow higher supply chain costs imposed due to increased order amplification to upstream. By considering the cost of upstream order amplification (or other relevant information cues in a different context), managers may be able to improve their policy decision and performance.

Acklowledgments

We are grateful to John Sterman, Benito Flores, Alan Love, Arzu Onar, two anonymous reviewers, and participants of the INFORMS, POMS, and System Dynamics Conference, for valuable feedback on earlier versions of this work, and to Lisa Noonan for her help running the games and capturing the data.

References

Adelman, D. 2001. First Loss Sets Gloomy Tone at CISCO. Wall Street Journal, May 9.

Armony, M., Plambeck, E., 2005. The Impact of Duplicate Orders on Demand Estimation and Capacity Investment. Management Science 51 (10), 1505-1518.

Bowman, E.H., 1963. Consistency and Optimality in Managerial Decision Making. Management Science 9, 310-321.

Chambers, J.M., Cleveland, W.S., Kleiner, B., Tukey, P.A., 1983. Graphical Methods for Data Analysis. Wadsworth Int. Group, Belmont, CA.

Chen, F., 1999. Decentralized Supply Chains Subject to Information Delays. Management Science 45 (8), 1076-1090.

Chen, F., Samroengraja, R., 1999. The Stationary Beer Game. Production and Operations Management 9 (1), 19-30.

Croson, R., Donohue, K., 2002. Experimental Economics and Supply Chain Management. Interfaces 32 (5), 74-82.

Croson, R., Donohue, K., 2003. Impact of POS Data Sharing on Supply Chain Management: An Experimental Study. Production and Operations Management 12 (1), 1-11.

Croson, R., Donohue, K., 2005. Upstream Versus Downstream Information and its Impact on the Bullwhip Effect. System Dynamics Review 21 (3), 249-260.

Croson, R., Donohue, K., 2006. Behavioral Causes of the Bullwhip Effect and the Observed Value of Inventory Information. Management Science 52 (3), 323-336.

Croson, R., Donohue, K., Katok, E., Sterman, J.D. 2005. Order Stability in Supply Chains: Coordination Risk and the Role of Coordination Stock. In Working Paper, Smeal College of Business, Penn State. State College, PA.

Diehl, E., Sterman, J.D., 1995. Effects of Feedback Complexity on Dynamic Decision Making. Organizational Behavior and Human Decision Processes 62 (2), 198-215.

Gonçalves, P., 2003. Phantom Orders and Demand Bubbles in Supply Chains. Ph.D. Dissertation, Sloan School of Management, MIT, Cambridge, MA.

Graves, S. 1999. A Single Item Inventory Model for a Nonstationary Demand Process. Manufacturing & Service Operations Management 1 (1), 50-61.

Greene, W.H., 1997. Econometric Analysis, 3rd ed. Prentice Hall, Upper Saddle River, NJ.

Gupta, S., Steckel, J.H., Banerji, A., 2001. Dynamic Decision Making in Marketing Channels: An Experimental Study of Cycle Time, Shared Information and Consumer Demand Patterns. In: A. Rapoport, and R. Zwick (Eds), Experimental Business Research. Kluwer Academic Publishers, Boston, MA., pp. 21-48.

Kaminsky, P., Simchi-Levi, D., 1998. A New Computerized Beer Game: A Tool for Teaching the Value of Integrated Supply Chain Management. In: H. Lee, and S.M. Ng (Eds), Supply Chain and Technology Management. Production and Operations Management Society, pp. 216-225.

Kampmann, C.E., 1992. Feedback Complexity and Market Adjustment: An Experimental Approach. PhD Dissertation, Sloan School of Management, MIT, Cambridge, MA.

Kleinmuntz, D., 1993. Misperceptions of the Implications of Feedback in Dynamic Decision Making. System Dynamics Review 9 (3), 223-237.

Lee, H., Padmanabhan, V., Seungjin, W., 1997a. Information Distortion in a Supply Chain: The Bullwhip Effect. Management Science 43 (4), 546-558.

Lee, H., Padmanabhan, V., Seungjin, W., 1997b. The Bullwhip Effect in Supply Chains. Sloan Management Review 38 (3), 93-102.

Mitchell, T.W., 1924. Competitive Illusion as a Cause of Business Cycles. Quarterly Journal of Economics 38 (4), 631-652.

Oehlert, G.W., 1992. A Note on the Delta Method. American Statistician 46, 27-29.

Oliva, R., 2003. Model Calibration as a Testing Strategy for System Dynamics Models. European Journal of Operational Research 151 (3), 552-568.

Pindyck, R.S., Rubinfeld, D.L., 1998. Econometric Models and Economic Forecast, 4th ed. Irwin, McGraw-Hill, Boston.

Plous, S., 1993. The Psychology of Judgment and Decision Making. McGraw-Hill, New York.

Richardson, G.P., Rohrbaugh, J.W., 1990. Decision Making in Dynamic Environments: Exploring Judgments in a System Dynamics Model-Based Game. In: K. Brocherding, O.I. Larchev, *et al.* (Eds), Contemporary Issues in Decision Making. Elsevier, Amsterdam, pp. 463-472.

Stata, 2003. Stata Release 8: Cross-Sectional Time-Series. Stata Press, College Station, TX.

Steckel, J.H., Gupta, S., Banerji, A., 2004. Supply Chain Decision Making: Will Shorter Cycle Times and Shared Point-of-Sale Information Necessarily Help? Management Science 50 (4), 458-464.

Sterman, J.D., 1987. Testing Behavioral Simulation Models by Direct Experiment. Management Science 33 (12), 1572-1592.

Sterman, J.D., 1989. Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. Management Science 35 (3), 321-339.

Sterman, J.D., 1992. Teaching Takes Off: Flight Simulators for Management Education. OR/MS Today 19 (5), 40-44.

Sterman, J.D., 2000. Business Dynamics: Systems Thinking and Modeling for a Complex World. McGraw-Hill, New York.

Tobin, J., 1958. Estimation of Relationships for Limited Dependent Variables. Econometrica 26, 24-36.

Tversky, A., Kahneman, D., 1974. Judgment Under Uncertainty: Heuristics and Biases. Science

185 (4157), 1124-1131.

Wooldridge, J.M., 2002. Econometric Analysis of Cross Section and Panel Data. MIT Press, Cambridge, MA.

Wu, D., Katok, E., 2006. Learning, Communication and the Bullwhip Effect. Journal of Operations Management (Forthcoming).

	Regressor	Model I	Model Ia	Model II	Model III	Model IV
β_L	Expected Loss (L_{t-1})	1.02 (0.01)***	1.02 (0.01)***	**1.00**	**1.00**	**1.00**
β_{LT}	Forecast Adjustment (G_{t-1})		0.02 (0.02)			
β_S	Inventory/Backlog (S_t)			-0.11 (0.01)***	-0.21 (0.01)***	-0.21 (0.01)***
β_B	Backlog ($S_t B_t$)				0.21 (0.02)***	0.21 (0.02)***
β_{SL}	Supply line (SL_t)			-0.01 (0.01)		0.00 (0.01)
β_0	Constant			0.33 (0.22)	1.69 (0.17)***	1.60 (0.25)***
	Log-likelihood value	-9769.9	-9769.4	-9535.8	-9426.0	-9425.8
	Wald χ^2	10395.0**	10391.8***	398.4***	562.5***	559.1***
	$R^{2(a)}$	0.54	0.54	0.58	0.60	0.60
	RMSE	4.83	4.83	4.60	4.46	4.46
	ρ	0.00 (0.00)	0.00 (0.00)	0.02 (0.01)*	0.04 (0.01)***	0.04 (0.01)***
	Backlog effect ($\beta_S + \beta_B$)				0.00	0.00
	P($\beta_{SL} + \beta_B$=0)				0.92	0.73
	Observations	3500	3500	3500	3500	3500
	Censored ($O_t \leq 0$)	520	520	520	520	520
	Number of players	100	100	100	100	100

Standard errors (SE) in parentheses: * significant at 10%; ** significant at 1%; *** significant at 0.1%.

(a) R^2=r^2, where *r* is the simple correlation between estimated and actual orders (Wooldridge, 2002).

† We found some evidence of autocorrelation in the residuals (p=0.08). Tests on the linear model (i.e., without the non-negativity constraint), however, indicated that, for all regressions, the underestimation on the SE resulting from the autocorrelation was not large enough to affect the reported significance of the estimates (in most cases the standard error of the estimates from the linear model matched that of the tobit model, and, when adjusting for autocorrelation in the residuals in the linear model, the change in the SE of the estimates was in the third significant digit). Since estimates with autocorrelation are unbiased, the reported results can be safely interpreted.

Table 1. Estimation results

	Regressor	Model III						Model IV		
		Full	R	W	D	F	~F	Full	~F	F
β_S	Inventory/Backlog (S_t)	-0.21 (0.01)***	-0.14 (0.01)***	-0.18 (0.02)***	-0.20 (0.02)***	-0.40 (0.03)***	-0.17 (0.01)***	-0.21 (0.01)***	-0.15 (0.01)***	-0.43 (0.03)***
β_B	Backlog ($S_t B_t$)	0.21 (0.02)***	0.15 (0.03)***	0.22 (0.03)***	0.19 (0.03)***	0.34 (0.05)***	0.17 (0.01)***	0.21 (0.02)***	0.18 (0.02)***	0.31 (0.05)***
β_{SL}	Supply line (SL_t)							0.00 (0.01)	0.02 (0.01)**	-0.09 (0.03)***
β_0	Constant	1.69 (0.17)***	1.30 (0.29)***	1.74 (0.28)***	1.60 (0.44)***	3.05 (0.57)***	1.44 (0.17)***	1.60 (0.25)***	0.77 (0.26)**	3.99 (0.67)***
	Log-likelihood value	-9426.0	-2161.7	-2367.5	-2455.3	-2128.1	-7193.6	-9425.8	-7188.4	-2123.3
	Wald χ^2	562.5***	110.7***	113.5***	139.5***	261.0***	349.8***	559.1***	348.2***	275.5***
	$R^{2(a)}$	0.60	0.16	0.43	0.49	0.77	0.43	0.60	0.43	0.77
	RMSE	4.46	2.94	3.88	5.33	5.11	4.18	4.46	4.17	5.09
	ρ	0.04 (0.01)**	0.07 (0.04)***	0.03 (0.02)	0.05 (0.02)***	0.10 (0.04)***	0.03 (0.01)***	0.04 (0.01)***	0.02 (0.01)***	0.12 (0.04)***
	Backlog effect	0.00	0.02	0.04	-0.01	-0.06	0.01	0.00	0.02	-0.12
	P($\hat{\beta}_S + \hat{\beta}_{SL}$ =0)	0.92	0.47	0.01	0.45	0.03	0.41	0.73	0.01	0.00
	Observations	3500	875	875	875	875	2625	3500	2625	875
	Censored ($O_t \leq 0$)	520	30	76	154	260	260	520	260	260
	Number of players	100	25	25	25	25	75	100	75	25

Standard errors in parentheses: * significant at 10%; ** significant at 1%; *** significant at 0.1%.
(a) $R^2=r^2$, where r is the simple correlation between estimated and actual orders (Wooldridge, 2002).
† See note on Table 1.

Table 2. Estimation results†

	Full	R	W	D	F	~F
$\bar{L}_{t-1}$	**7.59**	**7.56**	**7.68**	**7.78**	**7.76**	**7.67**
$\bar{\sigma}_{L_{t-1}}$	4.78	1.26	3.23	5.11	7.30	3.56
$S^* = \beta_0 / -\beta_S$ (†)	**7.94**	**9.44**	**9.63**	**8.12**	**7.56**	**8.57**
S.E.	0.67	1.65	1.21	1.91	1.19	0.83

(†) Calculations based on the "delta method" (Oehlert, 1992) values might differ from calculations based on coefficients from Table 2 because of rounding.

Table 3. Demand variance and estimated parameters–base model and by echelon

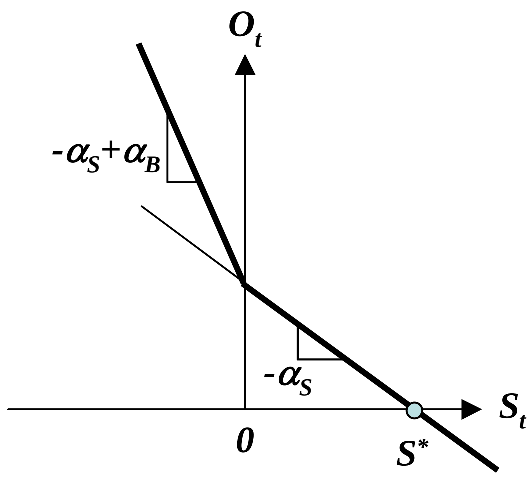

Figure 1. Order response to the inventory/backlog position

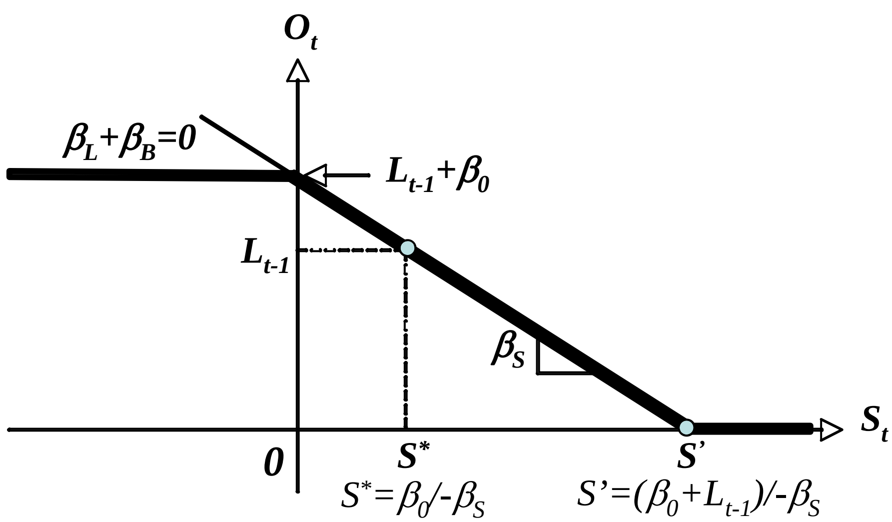

Figure 2. Estimated model's response to the inventory/backlog position

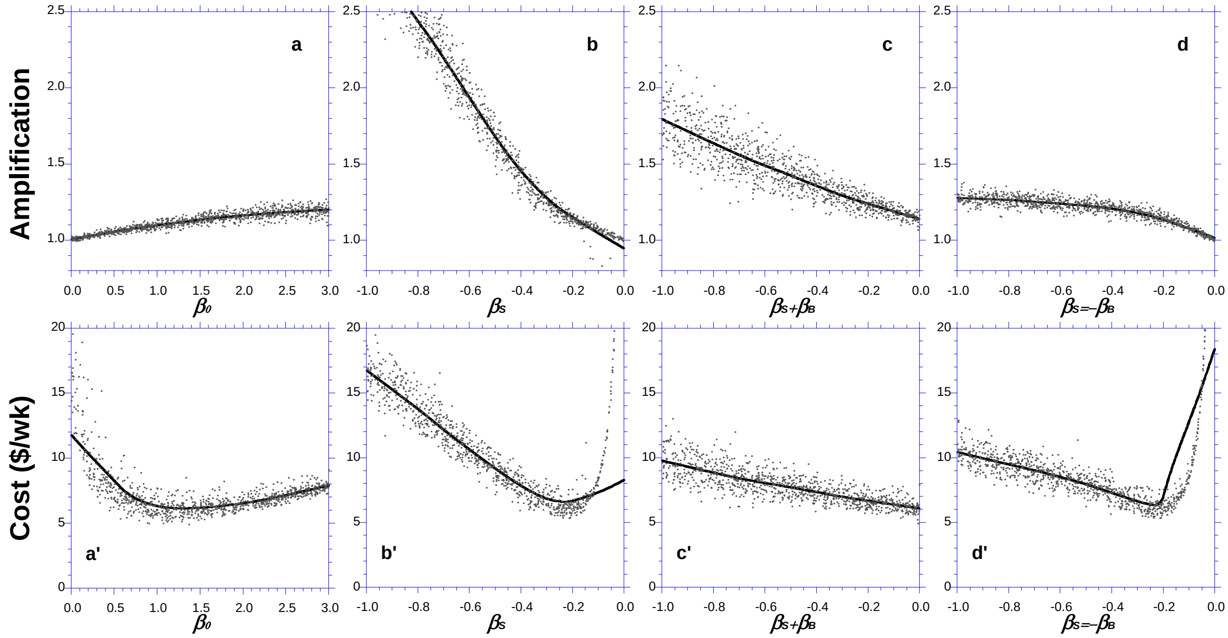

Smooth line estimated by means of the locally weighted regression scatter plot smoothing (lowess) procedure using, for each point, 40% of the sample (Chambers, Cleveland, Kleiner, and Tukey, 1983). The lowess line seems to deviate from the center of the data in the range (-0.21,0) in panels b and b' because the decision rule is unstable in that range and generated some extreme outliers. The deviation in panel d' is the result of extreme cost values when $\beta_S \rightarrow 0$.

Figure 3. Univariate sensitivity of decision parameters—Monte-Carlo simulations

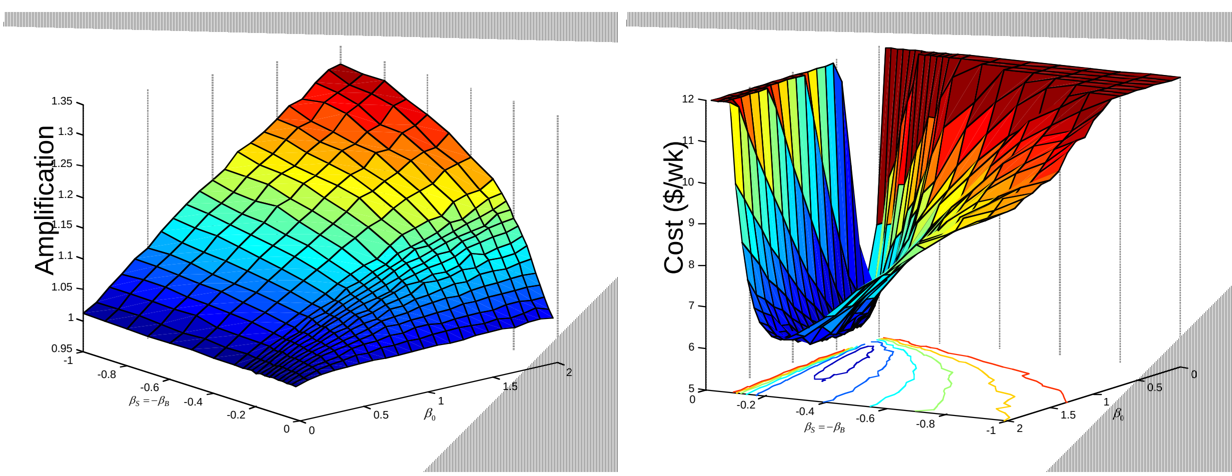

Note that x and y axes are reversed in the cost graph to give a better perspective of the response surface and that isocost curves are shown in the xy plane. Cost values have been truncated at 12 to show better detail of the minimal basin.

Figure 4. Multivariate sensitivity of decision parameters—Monte-Carlo simulations