




9-1997

Campbell Soup's Continuous Replenishment Program: Evaluation and Enhanced Inventory Decision Rules

Gérard P. Cachon
University of Pennsylvania

Marshall Fisher
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/marketing_papers

 Part of the [Business Administration, Management, and Operations Commons](#), [Business Analytics Commons](#), [Business and Corporate Communications Commons](#), [Management Information Systems Commons](#), [Marketing Commons](#), and the [Operations and Supply Chain Management Commons](#)

Recommended Citation

Cachon, G. P., & Fisher, M. (1997). Campbell Soup's Continuous Replenishment Program: Evaluation and Enhanced Inventory Decision Rules. *Production and Operations Management*, 6 (3), 266-276.
<http://dx.doi.org/10.1111/j.1937-5956.1997.tb00430.x>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/marketing_papers/255
For more information, please contact repository@pobox.upenn.edu.

Campbell Soup's Continuous Replenishment Program: Evaluation and Enhanced Inventory Decision Rules

Abstract

Campbell Soup's continuous replenishment (CR) program is a novel innovation designed to improve the efficiency of inventory management throughout the supply chain. With CR (1) retailers pay a constant wholesale price but continue to participate in consumer promotions, (2) retailers transmit to the supplier daily inventory information via electronic data interchange (EDI), and (3) the supplier assumes responsibility for managing retailer inventories, i.e., vendor managed inventories (VMI). We develop simple inventory management rules to operate CR, and we test these rules with a simulation using actual demand data provided by Campbell Soup. On this sample we find that retailer inventories were reduced on average by 66% while maintaining or increasing average fill rates. This improvement reduces a retailer's cost of goods sold by ~1.2%, which is significant in the low profit margin grocery industry. Furthermore, these savings could have been achieved without VMI.

Disciplines

Business | Business Administration, Management, and Operations | Business Analytics | Business and Corporate Communications | Management Information Systems | Marketing | Operations and Supply Chain Management

Campbell Soup's Continuous Replenishment Program:
Evaluation and Enhanced Inventory Decision Rules[□]

Gérard Cachon

Marshall Fisher

Fuqua School of Business

The Wharton School of Business

Duke University

University of Pennsylvania

Durham NC 27708

Philadelphia PA 19104

June 1995, revised September 1996.

[□] This research was presented at the Global Supply Chain and Technology Management Conference sponsored by the Hong Kong University of Science and Technology, June 1995. The authors would like to acknowledge the gracious assistance from the Campbell Soup Corporation, and in particular Ed James, Ron Ferner, and Harry Tetlow. Morris Cohen, Paul Kleindorfer, Hau Lee, Roy Shapiro, and Yu-Sheng Zheng provided many useful comments on earlier drafts. We are grateful for the helpful comments of the editors and the anonymous referees.

Abstract

Campbell Soup's Continuous Replenishment (CR) program is a novel innovation designed to improve the efficiency of inventory management throughout the supply chain. With CR: (1) retailers pay a constant wholesale price but continue to participate in consumer promotions; (2) retailers transmit to the supplier daily inventory information, via Electronic Data Interchange (EDI); and (3) the supplier assumes responsibility for managing retailer inventories, i.e. vendor managed inventories (VMI). We develop simple inventory management rules to operate CR, and we test these rules with a simulation using actual demand data provided by Campbell Soup. On this sample we find that retailer inventories were reduced on average by 66% while maintaining or increasing average fill rates. This improvement reduces a retailer's cost of goods sold by approximately 1.2%; which is significant in the low profit margin grocery industry. Furthermore, these savings could have been achieved without VMI.

The grocery industry has launched an ambitious effort called Efficient Consumer Response (ECR) to totally reengineer the supply channel. So far, the biggest reported benefits under ECR have come from Continuous Replenishment (CR), a new protocol for controlling the flow of information and product between a supplier and their retailer customers. Under CR retailers report daily to the supplier their recent demand and current inventory position. The supplier uses this information to determine replenishment shipments to each retailer. Since its inception in the grocery industry, Continuous Replenishment has been adopted in many other industries, often under the name Vendor Managed Inventory (VMI), to manage the flow of a wide range of daily consumables and basic apparel.

Campbell Soup is an early and successful adopter of Continuous Replenishment. Campbell Soup ships products from several plant warehouses to retailer distribution centers (DCs), from which the stores operated by the retailer are resupplied. The key features of the Campbell Soup Continuous Replenishment program are (1) products are sold to retailers at an Every-Day-Low-Price (EDLP), although retailers can run promotions to consumers, (2) every weekday morning retailers send to Campbell Soup via EDI their current inventory position and demand data from their DCs, (3) based on this information Campbell Soup sends a resupply shipment to the retailer DC. Each day a minimum inventory level is determined for each product, and for any product below its minimum, a shipment

is made sufficient to bring the product inventory to at least the minimum level. Only full truckloads are shipped and additional volume is added to the order as necessary to bring the total shipment to an integral number of truckloads.

In early tests with four retailers Campbell Soup found that Continuous Replenishment reduced inventories in retailer DCs by 50%; while increasing service slightly from 98.7% to 99.5%. The impact on Campbell Soup's own inventory was not tracked because it was felt that the percentage of volume on Continuous Replenishment at the early stage was insufficient to expect significant impact on Campbell Soup's operations. Nonetheless, it was felt that demand volatility for Campbell Soup had decreased somewhat due to the additional volume on EDLP.

In the Spring of 1993 we launched a project sponsored by Campbell Soup through the Wharton Fishman-Davidson Center for Service and Operations Research Partnership. Campbell Soup met with us to describe the environment within which Continuous Replenishment operates. In addition, they provided one year's worth of data for several hundred SKUs at four retailers. The data include demands at the retailers' DCs as well as well as the shipments that had been made during this period by Campbell Soup's Continuous Replenishment algorithm. The question was then whether more sophisticated inventory rules could improve on performance.

We developed a forecasting and inventory management system with the fol-

lowing features: An exponential smoothing model determines forecasts during non-promotion periods; the forecasts are used to determine for each SKU reorder levels that are appropriated given the desired service level, the SKU's average demand, the SKU's demand volatility and the retailer's average demand rate relative to truck capacity; orders are packed into full truckloads in decreasing priority order; and demand during past promotions are used to forecast future promotion demand.

This algorithm reduced inventories by 66% relative to what the retailers had been carrying and by 33% relative to what Campbell Soup had achieved. Thus this paper provides further confirmation of the benefits achievable through Continuous Replenishment and presents an algorithm that is highly effective relative to current industrial standards. We also determine how inventory varies with targeted service level and find that inventory requirements rise sharply if we target service above 99.5%

1. Data description and operating assumptions

To conduct this study Campbell Soup provided us with data from the distribution centers of four retailers as well as a description of their operating environment. The latter includes information regarding typical order processing and transportation times, minimum order quantities and promotion characteristics.

Over a one year period the data include for each retailer, each day and each product (1) in stock inventory, (2) open interest, (3) withdrawals, and (4) cuts. A product is a distinct stock keeping unit (SKU), so 8 oz. and 12 oz. chicken noodle soups are considered different products. In stock inventory only includes inventory in the retailer's distribution center, and does not include any store level inventory. Open interest equals the number of cases that have already been scheduled to be shipped to the retailer, but have not arrived. Withdrawals are the actual number of cases shipped out of the retailer's distribution center on the previous day. Cuts are the number of cases requested by the retailer's stores but were not shipped because the product was not available in the retailer's distribution center.¹ Hence, demand out of a retailer's distribution center equals withdrawals plus cuts. All demands are eventually filled. For each of the four distribution centers, Figure 1 displays weekly demand for all SKUs. The retailers experience substantially different total demand, but their seasonal patterns are similar.

In addition to the significant amount of yearly seasonality in demand, there is a substantial amount of demand "seasonality" within each week. On average across the four retailers Tuesday demand represents 22% of total demand, the

¹Cuts also include damaged products that cannot be sold. Campbell Soup does not record damaged products separately, but they indicated that damaged products constituted a small portion of total cuts.

remaining weekdays average about 15.5% of total demand and the weekend represents about 16% of total demand. However, the pattern of weekly seasonality differed somewhat across the retailers.

The data include 334 distinct SKUs, but each retailer carried between 163 and 242 different SKUs. Furthermore, some of the SKUs are not sold throughout the year either because they are seasonal or they are new products. Forecasting demand for these SKUs is particularly difficult, so they were not included in the study. (The SKUs dropped from the study represented less than 13% of the total demand.) Overall between 128 and 171 SKUs per retailer were included. Of these remaining SKUs about 40% are carried by all four retailers and almost 60% are carried by three or more retailers.

1.1. Shipping parameters

Campbell Soup estimates that if it decides to ship inventory to a distribution center on day d , then the inventory is available to be withdrawn from the distribution center on day $d + 2$: Past experience indicates that Campbell Soup generally has sufficient stock in its warehouse to fill any retailer need, no matter what quantity is requested. In addition, CR customers are given some priority, so whenever inventory is tight at Campbell Soup, the CR customers are most likely to be served.

Shipments of inventory to retailers are always constructed in full truck load quantities, which equals approximately 20 pallets. The number of cases per pallet, Q_{is} , varies by product, but most products are either 60 or 120 cases per pallet.² It is assumed that each product is always shipped in full pallet quantities, and the Campbell Soup facility serving each distribution center stocks all of the SKUs. Hence, a full truck can be loaded with up to 20 different SKUs and all SKUs are available to be loaded onto a truck. Assuming full truck load quantities and full pallet product ordering ensures that any observed inventory savings will not come at the expense of increased transportation or handling costs.

1.2. Consumer promotions

Retailers that adopt CR purchase product at a constant wholesale price, but they retain full control over pricing to consumers. In fact, CR retailers continue to engage in a variety of merchandising promotions which include “end-aisle displays”, coupons, “loss leaders” and simple price discounts.³ Implementing a promotion generally requires the retailer to withdraw large quantities of product from its distribution center before the promotion is advertised to the public. Figure 2

²The actual number of cases per pallet for each product is assumed to equal the modal minimum order size determined from the data. In most cases this minimum order quantity represents over 90% of orders for the product.

³An end aisle display is a placement of a large amount of product at the end of a aisle in a grocery store. Since many shoppers travel around these points as they navigate the store, this location can serve to increase a product’s sales. Loss leaders are products sold at a deep discount to attract consumers to the store.

displays the demand pattern for three products which apparently experience consumer promotions, and one product without promotions for comparison. These time series indicate that promotions tend to be quick surges in demand: They often last less than a week and often occur over only one day.

Campbell Soup indicated that under normal operations, the timing of promotions is generally known in advance of the actual promotion, because retailers conduct promotions only after receiving some inducements from Campbell Soup. Since the data we received do not identify the occurrence of a promotion, for the purpose of our study, this information must be estimated from the demand data.

Table 1 presents the procedure used to identify consumer promotions. Steps 1-3 are designed to determine measures of normal demand and the variability of demand. Medians are used to estimate “average demand” because the presence of promotions tends to upwardly bias arithmetic averages. Step 4 identifies promotions as a series of days with demand extremely high relative to normal demand. Multiple conditions are used because some promotions are 1 day “extreme spikes” in demand whereas other promotions are “very high demand” over several consecutive days. In the application of these criteria zero demand days are ignored. Thus, a promotion with four demand days may actually cover more than four calendar days. Step 4 also includes conditions to ensure that a promotion represents a significant increase in demand for a particular day of the week. This

is important for products that exhibit strong weekly seasonality. Step 5 indicates that the first three steps are repeated to ensure that promotions found in the first iteration do not bias the assessment of true normal demand. Finally, since some promotions occur from one weekend to the next, promotions in close time proximity are combined. This procedure for identifying promotions requires assumptions, but short of actual data there does not appear to be a purely objective measure of identifying promotions.

1.3. Promotion characteristics

Several interesting observations follow from the analysis of promotions. Approximately half of all SKUs in the study have at least one promotion during the year but over 90% of SKUs have 2 or fewer promotions during the year. Promotions tend to be quick. About three quarters of the promotions last three days or less and about 95% of promotions are over within seven days. About 60% of promotions start on Tuesday or Wednesdays. Demand during many promotions is often dramatically greater than median daily demand: Demand in 54% of promotions is more than 15 standard deviations greater than daily median demand, and demand in 3% of promotions is more than 100 standard deviations greater than normal daily demand. However, promotion demand represents a relatively small percentage of total yearly demand for most products. For 90% of products

promotion demand is less than 15% of total yearly demand, and promotion demand is less than 20% of total demand among 90% of the products with one or more promotions.

1.4. Inventory objectives

Campbell Soup attempts to minimize inventory while providing each retailer and each product a 99% fill rate both during normal and promotion demand periods. Furthermore, shipments to retailers should equal a full truck load and shipments of each product must equal an integer multiple of pallets.

2. Enhanced Inventory Management

This section describes several inventory management rules that daily decide how much of each SKU to ship to each retailer. The objective is to minimize the retailer's inventory while maintaining a 99% fill rate for each SKU.

To test our approach we divide the data into two periods. The "calibration period" is the first 182 days and the "test period" is days 183-365. All of the parameters used in the inventory management rules are estimated with calibration period data. All of the results regarding inventory or service are derived from the test period data.

A comment on notation, for each variable defined, unless otherwise noted, an

"i" in the subscript refers to retailer $i \in [1; 4]$, a "s" in the subscript refers to SKU $s \in [1; 171]$; and a "d" superscript refers to day $d \in [1; 365]$.

2.1. Forecasting normal demand

Define several variables for day d : $p_{is}^d = 1$ indicates that a promotion occurs on day d , otherwise $p_{is}^d = 0$; when there is no promotion D_{is}^d equals actual demand, otherwise $D_{is}^d = 0$; when there is a promotion PD_{is}^d equals actual demand, otherwise $PD_{is}^d = 0$. Let F_{is}^d equal the forecast of normal demand (i.e.e, non-promotion demand) for day d , where each F_{is}^d is evaluated using a simple exponential smoothing model:

$$F_{is}^d = \begin{cases} F_{is}^{d-1} + \alpha (D_{is}^{d-1} - F_{is}^{d-1}) & \text{if } p_{is}^{d-1} = 0 \\ F_{is}^{d-1} & \text{if } p_{is}^{d-1} = 1 \text{ and } F_{is}^{d-1} > 0 \\ D_{is}^{d-1} & \text{otherwise} \end{cases} \quad (2.1)$$

where α is a constant. This forecasting rule makes several assumptions about the demand process: promotion demands have little effect on subsequent normal demand because a forecast is not updated if it occurred on a promotion day; and demand on a particular day of the week depends only on the past history of demand on the same day of the week. Thus, demands on Mondays have no effect on Tuesday demands. Finally, it is assumed that a single constant α can

effectively apply to all retailers and all SKUs.

The parameter τ is chosen to minimize a measure of forecast errors in the calibration period. For each retailer, day and SKU combination let the forecast error equal $D_{iS}^d - F_{iS}^d$, assuming no promotion occurs on day d . For each $\tau \in [0:01; 0:99]$ in 0:01 increments the sum of squared forecast errors for all retailers and all products in the calibration period was calculated. The parameter setting $\tau = 0:31$ minimized this value and hence it is used in all subsequent forecasts.

2.2. Forecasting promotion demand

With promotions it is particularly difficult to balance the trade-off between providing good service and carrying low stocks of inventory. If a promotion forecast is low then substantial shortages occur. But if the promotion forecast is high, the retailer may be left at the end of the promotion with a supply of inventory sufficient to fill several months worth of normal demand. Unfortunately, demand before a promotion provides little warning regarding the level of promotion demand. For example, demand in the first day of a promotion is at least five times greater than the average demand in the five days before the promotion for 43% of promotions, and for 5% of promotions the first day's demand is more than twenty times average pre-promotion demand. Hence, it is apparent that merely observing daily withdrawals from a retailer's distribution center will provide a

poor system for managing promotion demand. For example, if Campbell Soup were not informed about pending consumer promotions, Campbell Soup would have to maintain in the retailers' distribution centers inventory equal to 70 times average daily demand to ensure that 99% of promotion demand is covered.

With a CR system each retailer must inform the supplier when a promotion begins, when it ends, and a forecast for the size of the promotion. Although the timing of the promotions are not recorded in the data set, they were known to Campbell Soup with certainty. While Campbell Soup used these forecasts to make its inventory decisions prior to each promotion, they were not recorded and hence are absent from the data. Therefore, we develop our own forecasting rules.

We forecast promotions using three pieces of information. First let CF_{is}^d equal the maximum inventory position at the retailer distribution center in the Δ days before a promotion that begins on day d . This is taken as a proxy for the forecast the retailer provided to Campbell Soup since this forecast was presumably used by Campbell Soup to decide how much inventory to ship to the retailer in advance of the promotion. Second, since products with higher demand may have higher promotions, let M_{is} equal median demand for each product. Finally, define

$$PP_{is}^d = \begin{cases} \sum_{t=1}^{\Delta} n_{is}^d \cdot \frac{1}{\Delta} \sum_{t=1}^{\Delta} PD_{is}^t = \frac{1}{\Delta} \sum_{t=1}^{\Delta} p_{is}^t & \text{if } p_{is}^d = 0; p_{is}^d = 1; \frac{1}{\Delta} \sum_{t=1}^{\Delta} p_{is}^t > 0 \\ 0 & \text{otherwise} \end{cases}; \quad (2.2)$$

where n_{is}^d is the number of days in a promotion that begins on day d . Note that for each product PP_{is}^d is defined only after the first promotion has been completed.

Ordinary least squares regression is used to determine PF_{is}^d , the forecast of total demand during the promotion that begins on day d :

$$\ln(PF_{is}^d) = \begin{cases} a_0 + b_0 \ln(M_{is}) + c_0 \ln(CF_{is}^d) & PP_{is}^d = 0 \\ a_1 + b_1 \ln(M_{is}) + c_1 \ln(CF_{is}^d) + d_1 \ln(PP_{is}^d) & PP_{is}^d > 0 \end{cases} \quad (2.3)$$

(Natural logs are taken to reduce heteroskedasticity.) If no promotion begins on day d then $PF_{is}^d = 0$. Using only the promotions in the calibration period, the following displays the parameters selected for both the first promotion for a product and the subsequent promotions (promotions 2 and higher)⁴:

	a	b	c	d	R ²	σ_p	σ_p^3
First promotion:	0.27	0.55	0.53	0.00	0.56	0.04	0.46
After first promotion:	0.18	0.48	0.25	0.37	0.83	0.01	0.39

In the above σ_p and σ_p^3 are the mean and standard deviations of the error terms respectively.

⁴The parameters for the “after the first” promotion parameters exclude retailer 2’s promotion on day 153 for UPC 105. This outlier significantly deteriorated the quality of the forecast in subsequent tests. It is assumed that Campbell Soup would have been able to identify this single promotion as an outlier. The parameters a ; b ; c and d for the first promotion are determined without including any promotion from retailer 4. Several of retailer 4’s initial promotions were dramatically different than forecasts. It is assumed that these errors were due to inexperience with the program since retailer 4’s forecasts improve dramatically with time.

The value PF_{iS}^d is a forecast for expected total demand during a promotion, but our objective is to ensure a 99% promotion fill rate. We define PR_{iS}^d as the stock required at the retailer's distribution center at the start of a promotion to ensure that 99% of the total promotion demand is filled assuming no additional replenishment arrives during the promotion. Letting $\Phi(R)$ and $\phi(x)$ equal the distribution and density functions respectively of the standard normal distribution,

$$PR_{iS}^d = PF_{iS}^d + e^{1/\rho + R_z^3/\phi};$$

where R_z is a constant implicitly defined by

$$0.99 = \Phi(R_z) + \int_{R_z}^{\infty} e^{3/\rho(R_z - x)} \phi(x) dx: \quad (2.4)$$

2.3. Ordering inventory

Each day for each retailer and each product a reorder point, R_{iS}^d ; is calculated. Comparing R_{iS}^d with a product's inventory position IP_{iS}^d (on hand inventory plus on route inventory minus backorders) determines the number of pallets required to raise the product's inventory position above R_{iS}^d . The b^h pallet in this set is assigned priority $R_{iS}^d - IP_{iS}^d + (b_i - 1)Q_{iS}$. For each retailer, pallets are loaded onto trucks in decreasing priority order. A truck is shipped only if there is a sufficient

number of pallets to fill the truck, hence some of the lowest priority pallets may not be shipped.

The actual reorder point for a product is calculated with the following formula:

$$R_{is}^d = \sum_{t=d}^{d+X+\phi_i} (F_{is}^t + PR_{is}^t) + (M_{i1} + M_{i2}) \frac{3}{4} \sigma_{is} \sqrt{2 + \phi_i}; \quad (2.5)$$

where $\frac{3}{4} \sigma_{is}$ is the standard deviation of forecast errors in the calibration period, and ϕ_i , M_{i1} and M_{i2} are positive real constants. (Note that ordering decisions on day d must occur before observing day d demand.)

In (2.5) the summation reflects the forecast for normal and promotion demand in the near term. The final term in (2.5) accounts for the variation in normal demand. The constant ϕ_i is included because the full truckload requirement implies that some low priority pallets may experience a shipping delay. It is likely that ϕ_i is highest for low demand retailers because they take longer on average to fill a truck load. Finally, the constant M_{i1} and M_{i2} are included to take into consideration changes in the variability of demand. Specifically, we observed that for products with normal demand and no promotions the variability in demand tends to be less than products with promotions. Even with promoted products, demand variability in the month after a promotion tends to be greater than demand variability before the promotion. Thus, $M_{i1} > 0$ and $M_{i2} = 0$ during

normal demand in which there has not been a promotion in the last month and $M_{i1} = 0$ and $M_{i2} > 0$ during the month after any promotion.⁵

Using (2.5) to set reorder points contains several limitations. For example, Q_{is} is not considered when setting a reorder point, even though all things being equal, the same fill rate can be maintained with lower reorder points for products with larger Q_{is} .⁶ In addition, this rule assumes that it is always too costly to ship in less than full truck load shipments. Another truck filling rule could state that a truck will be shipped when a certain cumulative priority level has been achieved by the loaded batches, whether the load fills the truck or not.

Determining the average inventory and service for any particular policy (M_{i1} , M_{i2} , ϕ_i) is difficult. Instead, a simpler approach is taken to select a policy. Using the calibration period data, average inventory and service are measured for each of the policies in the set containing all combinations of the following parameters: $M_{i1} \in \{0, 0.25, \dots, 8\}$; $M_{i2} \in \{0, 0.25, \dots, 8\}$; $\phi_i \in \{1, 2, \dots, 6\}$. The chosen policy was the one that minimized average inventory while providing at least a 99.5% fill rate:

⁵There are other techniques to cope with changing variability. For example, forecast errors could be estimated for both normal demand without a recent promotion and normal demand with a recent promotion. However, if a product had not experienced a promotion in the calibration period then there would be no estimate of post promotion volatility for the test period. In this case, the additional volatility would have to be inferred from the pattern of other products.

⁶Larger Q_{is} reduces the frequency of orders, and thus reduces the proportion of time the product is exposed to stockouts.

	M_1	M_2	ϕ_i
Retailer One:	3	4	2
Retailer Two:	3	4	2
Retailer Three:	1.5	4	3
Retailer Four:	1.5	4	4

3. Results

Several results are presented on the performance of our inventory management rules on the test period data. Recall that these rules are applied to these data as if they were being used to conduct actual operations.

Table 2 displays some results from the promotion forecasting model. In the table “excess inventory” equals the total amount of inventory at the retailer’s distribution center at the end of the promotion. Ideally, the fill rate during the promotion would be high and the amount of excess inventory would be low. Table 2 demonstrates that the promotion rule provides good results in the test period even though it is based exclusively on parameters estimated with calibration period promotions. Furthermore, our rules lower promotion related inventory while maintaining good service (CR’s average fill rate equals about 97.5% during promotions).

Table 3 presents results for the test period including both the promotion

forecasting and normal demand inventory management. In the table “Before CR” is an estimate of the retailers’ performance managing their own inventories before joining the CR program. Campbell Soup estimated that they eliminated approximately 50% of these retailer’s inventory, hence the Before CR turns are approximately half of the CR turns.

Overall, our inventory management rules perform quite well even relative to CR both in terms of inventory and service. Average inventory, compared to CR dropped between 20% and 61%. (Average inventory for each retailer equals the average across all products of each product’s average inventory.) Compared with the Before CR results, our rules reduce inventory 66% on average. This translates into a savings of about 1.2% of the cost of the product, assuming inventory carrying charges equal to 25% of the cost of inventory. Given that profit margins in the grocery industry are usually less than 3%, this cost savings is substantial.

Fill rates for these policies ranged from 99.0% to 99.9%. However, in some cases the fill rate level with our rules was slightly lower than that obtained by CR. Hence, some of the inventory reduction may be due to lower fill rates rather than better inventory management.

Our rules were set to target at least a 99% service level. However, the rules can be used to determine the amount of inventory that would be carried at different

service levels, which is displayed in Figure 3. The figure indicates that while inventory levels increase with the desired service, as expected, the increase is relatively minimal until service is set at about 99.5%. This was a useful insight to the Campbell Soup, because some people in the company had believed that a substantial inventory savings could be achieved by lowering the targeted service level to 95%.

4. Discussion

We conclude that simple inventory management rules can significantly reduce inventories in the retailers' distribution centers. These rules are effective because (1) each product's individual characteristics, such as average demand and demand variability, are incorporated into determining replenishment decisions, (2) promotion forecasting is enhanced by considering past promotions, and (3) the full truck load constraint is not ignored. While additional inventory savings could probably be achieved if the retailers implemented fewer consumer promotions, this result indicates that significant efficiency gains nevertheless can occur despite the presence of consumer promotions.

Our test was conducted in the soup industry, but we feel the results apply to some other industries. The following lists the important characteristics of the soup industry: Relatively stable set of SKUs; seasonal demand; extensive consumer

promotion activity; and low profit margins. A stable product set suggests that past sales are a valuable source for short term forecasting, but demand surges due to consumer promotions indicates that these events must be handled specially. Low profit margins implies that inventory reductions should not come at the expense of increased handling and transportation costs, hence the use of minimum order quantities and full truck load shipping. These characteristics apply to most dry goods in the grocery industry, consumer hardware supplies, perennial selling toys and many goods sold through general merchandisers.

Interestingly, the our rules do not need to be implemented by Campbell Soup. Instead, the retailers could implement these rules and merely transmit to Campbell Soup each day their orders, and the same results would have been obtained. Hence, in this study the inventory reductions cannot be attributed to vendor management. However, we do not wish to conclude that vendor management is never important. It is possible that demand and inventory information would allow Campbell Soup to better allocate scarce inventory among retailers or this information could improve Campbell Soup's production scheduling. Indeed, Clark and Hammond (1995) provide empirical evidence indicating that only EDI ordering does not lower retailer inventories as much as EDI ordering combined with vendor management. Clearly, additional theoretical and empirical work is required to resolve the benefits directly attributable to vendor management.

Table 1: Procedure for identifying consumer promotions

Step 1	For each product s at each retailer evaluate (1) the median daily demand across the entire year, M_{is} , and (2) the median daily demand across the entire year for a particular day of the week, M_{is}^w .
Step 2	Evaluate two standard deviations of demand: $\sigma_{is} = \sqrt{\frac{1}{n_i - 1} \sum_{d=1}^{365} (D_{is}^d - M_{is})^2}$ $\sigma_{is}^w = \sqrt{\frac{1}{n_i^w - 1} \sum_{d=1}^{365} (D_{is}^d - M_{is}^w)^2}$ where D_{is}^d is consumer demand on day d , n_i equals the number of days with positive demands, and n_i^w equals the number of days with positive demand that are the w^{th} day of the week.
Step 3	Assign each day a "normalized demand" : $Z_{is}^d = \frac{(D_{is}^d - M_{is})}{\sigma_{is}}$ $Z_{is}^{dw} = \frac{(D_{is}^d - M_{is}^w)}{\sigma_{is}^w}$
Step 4	A series of days passing any one of the following three conditions is labeled a promotion: A) One or more days with $Z_{is}^d \geq 5$, and $Z_{is}^{dw} \geq 3$: B) Two or more consecutive positive demand days with $Z_{is}^d \geq 4$, and $Z_{is}^{dw} \geq 3$: C) Four or more consecutive positive demand days with $Z_{is}^d \geq 3$, and $Z_{is}^{dw} \geq 3$:
Step 5	Repeat steps 1-3 but exclude any days labeled as a promotion in the first iteration in the evaluation of M_{is} and M_{is}^w :
Step 6	Combine any promotions within 5 or less days of each other.

Table 2: Performance of the promotion forecasting model:

	Promotion type	Our promotion forecasting model excess inventory	CR's excess inventory	% change	Our model's service
Calibration period	First	17,467	50,841	-65.6%	98.9%
	After the first	52,887	72,380	-26.9%	95.8%
Test period:	First	42,662	86,179	-50.5%	98.7%
	After the first	34,839	45,252	-23.0%	99.6%

Table 3: Change in inventory turns and service across retailers (test period results)

Retailer	Total yearly demand	Inventory turns			Service	
		Before CR	CR	Our model	CR	Our model
1	931,755	12.7	25.5	34.7	99.5%	99.0%
2	1,444,244	19.1	38.1	53.2	99.7%	99.9%
3	471,815	7.9	15.9	28.4	99.5%	99.3%
4	316,072	12.6	25.2	33.0	99.6%	99.7%

Figure 1: Weekly demand at each retailer distribution center for all SKUs




Figure 2: Daily retailer demand for four selected products.


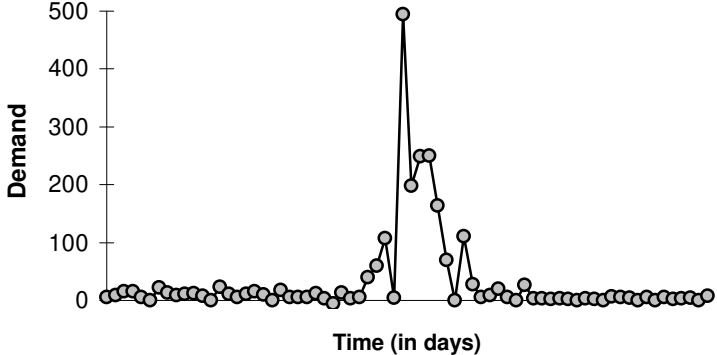




Figure 3: Inventory and service frontier at each retailer using our inventory management model

