The Value of Information for Managing Inventory in Vending Machines

Michael E. Ketzenberg
MKetzenberg@tamu.edu
Mays School of Business
Texas A&M University
College Station, TX, USA
979.845.1616

Neil Geismar
Mays School of Business
Texas A&M University
College Station, TX, USA
NGeismar@mays.tamu.edu
979.458.0033

Richard D. Metters
Mays School of Business
Texas A&M University
College Station, TX, USA
979.845.1616
RMetters@mays.tamu.edu
979.845.1616

Erwin A. van der Laan
Rotterdam School of Management, Erasmus University
Rotterdam, The Netherlands
elaan@rsm.nl
(31) (0) 10 44082763 / 2032
Fax (31) (0) 10 44089010

The Value of Information for Managing Inventory in Vending Machines

---

An important difference between both manufacturing and wholesaling versus retail is the information available concerning inventory. Typically, far less information characterizes retail. Here, an extreme environment of information shortfall is examined. The environment is technically termed "unattended points of sale," but colloquially called vending machines. Once inventory is loaded into a machine, information on demand and inventory level is not observed until the scheduled reloading date.

Technological advances and business process changes have drawn attention to the value of information in retail inventory in many venues. Moreover, technology is now available that allows unattended points of sale to report inventory information. Capturing the value of this information requires changes in current business practice. We demonstrate the value of capturing information analytically in an environment with restrictive demand assumptions. Experiments in an environment with realistic demand assumptions and parameter values show that the value of information depends greatly on operating characteristics and can range from negligible effects to increasing profitability 30% or more in actual practice.

Keywords: retail, value of information, vending machines, inventory management

1 Introduction

Traditional economic theory states that information and inventory are substitutes: more information reduces the need for inventory (e.g., Milgrom and Roberts 1988). In general, the value of information (VOI) literature in operations management seeks to test this theory and to determine the conditions under which it is true. Various researchers have determined the VOI for different types of information in a wide variety of operating environments. We seek to extend the literature to the case in which there is a lack of knowledge of inventory level in current practice. Complete inventory level knowledge is not the case in any environment, but it is closer to being true in some arenas than others. In the spectrum of inventory level knowledge, the more controlled environments that are found in the manufacturing and warehousing sectors provide relatively accurate knowledge of inventory level. In retail, however, this information is more difficult to obtain. The products being inventoried are in individual units stocked by hand rather than carton loads carried by forklifts, and customers, rather than employees, have more direct control over the inventory. In one study, a *majority* of retail inventory records were found to be inaccurate (DeHoratius and Raman 2008).

We study a different environment than DeHoratius and Raman, one that is at the other end of the spectrum of inventory level knowledge from a tightly controlled warehouse. Rather than the known, but inaccurate, numbers in a typical retail store computer system, we study an environment in which the inventory level is unknown by design. Our general setting here is “unattended points of sale.” Broadly, this work could be applied to many situations in which employees do not continuously track inventory status. Unattended points of sale with stochastic demand also occur in customer-held inventory environments such as propane/oil delivery (Bell, et al. 1983, Dror and Ball 1987), garbage removal (Beltrami and Bodin 1973), and recycling pickup.

The specific environment of our study calls itself “automatic merchandising,” but it is colloquially referred to as “vending machines.” In industry practice, a machine is loaded with inventory and then is essentially abandoned until the next scheduled loading time. No sales or inventory information is sent back during the intervening period, and no employees are even able to perform a visual check since the machines are not at the replenishing firm’s location. However, technology is now available that can provide inventory level information for remote sites. We seek to understand and evaluate the potential value that having this information may yield.

Our analytic and experimental results conclude that the VOI of inventory level knowledge is not as straightforward as economic theory suggests. For certain parameter settings, the benefits of information are low, with profitability increases of less than 1%. Other parameter values indicate that new technologies that increase the information flow can be very beneficial, leading to profitability increases of more than 28%. By extension, other retail environments with low information content may also be positively affected. However, capturing the VOI must be accompanied by changes in business practice.

After a description of the environment, this work will be positioned with respect to the literature. An analytic model will be formulated to demonstrate VOI under restrictive demand assumptions. Experiments will then show the VOI in a broad array of realistic business situations.

2 The Vending Environment

Vending machines are ubiquitous, have significant economic impact, and can benefit from advances in research. The industry trade association, the National Automatic

Merchandising Association (NAMA), states that there are 5.4 million beverage, snack, and food vending machines in the U.S., garnering $20 billion in annual sales (Maras 2010). Japan may be the world's per capita leader in vending machine consumption, with 5.6 million machines generating $58 billion in annual revenue (Ito 2007).

In the last several years, there have been rapid advancement in information technologies to improve business performance. Technology is now available to monitor machines remotely. Machines can "call in" via the internet to report inventory level. This technology is relatively inexpensive, costing approximately $150/machine to install. There is also a monthly telephony charge of $5-$10 for the lead machine in an area (the machines are networked and only one machine is needed to transmit data).

Despite the potential for better information, vending companies have been slow to adopt this new technology. Essentially, while the cost of these technologies is easy to determine from the marketplace, information on their benefits and returns on investment are scarce. Hence, there is a real need for research on the general question of the trade-off between information and inventory, as well as on the specific industry.

The industry in the U.S. has $20B in annual sales, but it is spread over more than 9,000 firms (Maras 2010, p.33). The prototypical firm is small, with roughly $2M in annual revenue (Maras 2010, p.33). The two firms under study were of this type, each with approximately 500 vending machines and half a dozen trucks. We obtained detailed data by job shadowing workers to gather time and motion information from one firm and by analyzing detailed sales figures from the other.

We now provide an overview of the replenishment operations—both with and without the new technology—then provide an illustration of the benefits that a company could realize from the new technology.

2.1 Operational Details

A traditional snack vending machine holds approximately 40 products, whereas a cold beverage machine may contain 8 to 20 products. The number of machines at a client site can range from several to a few hundred. In current practice, groups of machines are visited and re-stocked on a cycle. This cycle may be once a week or "every Tuesday and Thursday," but for sake of notational convenience, we describe this mathematically as a cycle of visiting once every R days. The practical range of R is one to 14. Each worker is typically responsible for 50 to 120 machines (NAMA 2008). The inventory involved is physically small (e.g., candy bars, soda cans), so the relatively small trucks employed have far more capacity for inventory than needed for a work shift. Hence, the main constraint is the employee time required to visit a machine, assess its needs, and restock. In the vending machines that are the subject of this research, there is very little deterioration or disposal of inventory, as the shelf lives of the products exceed the review periods by orders of magnitude. Truck routing is also not significant because the daily route for a truck is often confined to an industrial park or other small area, so only a small percentage of the day is spent driving between sites and the order of site visits is not important. The bulk of the shift time is spent servicing machines, kitting inventory at the truck, and walking back and forth between the truck and the machines inside a client building.

In traditional inventory literature, the description above would lead to implementing a (Q,R) or (S,R) model. That is, given R time periods between inventory reviews, order either Q units or order-up-to S units of inventory, where the calculation of Q or S follows known procedures. The vending machine problem is a simplified version of the traditional (S,R) model in which $S = Y$ and R is the only decision variable. Practicalities militate against an order-up-to amount that is less than machine capacity. As will be discussed later, the standard deviation of demand is typically a multiple of the mean, so stocking less than capacity risks lost sales even for traditionally low mean demand products. Further, setting individual order-up-to targets for each of 40 products per machine, 100 machines per driver, where different products are frequently introduced, is operationally infeasible. It is estimated that 90% of the vending machine business is operated in this fashion – visiting on a cycle of R days and stocking to capacity (NAMA 2008).

Stocking to maximum capacity is not limited to vending machines. It has also been noted as standard practice in other retail applications (Ketzenberg, Metters and Vargas 2000, Borin, Farris and Freeland 1994, p.364). Marketing literature on retail shelf space allocation typically assumes all space is filled with product (e.g., Bultez and Naert 1988). Beyond operational simplicity, it is also believed that stocking to maximum capacity increases demand (Balakrishnan, Pangburn, and Stavrulaki 2004). The practice of stocking to capacity has been verified with the national association, several firms in the business, and by the authors personally shadowing route workers at their task.

The inventory literature speaks of "review" and "replenishment" as separate events. Vending machine inventory management is a special case in which review and replenishment are coincident events. On each machine visit, stock levels are evaluated and raised to machine capacity, i.e., a base-stock policy with zero lead time. Given that the next replenishment would take place in another R time periods, and that demand forecasting for the individual products in a machine is expensive, difficult, and highly error prone due to high levels of demand uncertainty, this is a reasonable managerial action.

As a consequence of the replenishment/review identity and the practice of filling the machine to capacity when it is replenished, the traditional decision in the vending environment is disarmingly simple: what should the replenishment period *R* be? If *R* is too small, too many expensive trips to the machine are made to no purpose. If *R* is too large, one incurs lost sales.

Accurate web-based remote machine monitoring complicates this decision-making by allowing the separation of review and replenishment. Rather than being reviewed and replenished every *R* periods, machines can now dynamically determine when service is needed. Essentially, this changes the inventory system from a (*R, Y*) system to a (*s, Y*) system, i.e., from periodic review to continuous review, where *Y* represents machine capacity. However, to take advantage of this ability, the typical practice of an operator running a fixed route must change. Dynamic scheduling of machine replenishment must occur.

An unusual aspect of the vending machine problem from an inventory theoretic perspective is the absence of holding costs. This is because, unlike in traditional commerce, the revenue from sales is kept in the machine, i.e., it is not usable by the machine's owner until replenishment occurs. Hence, a sale does not change the value of the contents of the machine, it merely exchanges goods for money. Therefore, the financial burden from holding inventory is not reduced upon sale. Further, an opposite exchange occurs when replenishment is made: machine inventory is brought back to capacity, the money is removed, and the value of the machine's contents remains the same. As a consequence, the cost of the inventory investment is independent of whether and when sales occur.

2.2 Illustration of Benefits of Smart Machines

We obtained two years' data for 55 snack machines that constitute a typical route for a vending machine operator in Texas. Under current practice, each machine is visited for review and replenishment on average every 3.6 days, with a range of 1-14 days. Based on cost and price information provided by the vending machine servicing firm, the total annual revenue generated by each machine is $3165 (range $904-$9944), providing an annual gross margin of $1813 (range $504-$5726). Hence, the total annual margin for all fifty-five machines is $99,810. Given the firms pro forma routing and labor costs, the average cost to visit a machine is approximately $6.50. Hence, the total annual cost for all replenishments is $36,387, which provides a total annual profit contribution of $63,423.

One convenient and simple method for using the new technology would be for the manager to set the visitation threshold to foster a 99% service level for 95% of the products in the machine. We calculated the effect this policy would have for the 55 machines in this route. Each machine's inventory is reviewed remotely at the beginning of each day. A machine will be visited for replenishment that day unless 95% of its products each has a 99% chance of meeting all demand over the next 24 hours – predicated on historical demand rates.

The preceding scheme was tested as follows. For each machine, based on the mean and standard deviation of the observed demand, we calculated the inventory level required to satisfy a day's demand with probability 0.99 for each product in that machine. If more than 5% of a machine's products are below their particular threshold levels, a replenishment visit is scheduled. This leads to an average replenishment interval of 5.0 days, which represents a reduction in the number of trips per year of 28.3% from current practice. This yields a savings of $10.44 per machine per month, or $6890 per year over all the machines in the route. In this case, the total annual profit contribution is $73,71, an increase of 16.2%.

3 Prior Literature

In retail, the true inventory level a firm faces can be difficult to discern for a number of reasons, including SKU misplacement by customers or employees, transaction or scanning errors, and other sources (Raman, DeHoratius and Ton 2001). In each of these situations, poor information can lead to a loss of sales. "Phantom stockouts" have been studied by Ton and Raman (forthcoming). These occur when items are in a store, but customers cannot locate them. They are related to the problem of "freezing" (DeHoratius, Mersereau and Schrage 2008), which occurs when there is no inventory that a customer can buy although the computer records indicate that inventory is available. Other representative work in inventory inaccuracy includes modeling the inventory accuracy benefits from RFID (de Kok, Van Donselaar, and Van Woensel 2008, Lee and Ozer 2007, Kang and Gershwin 2005) and balancing increased inventory inspection with the benefits of inventory accuracy (Kok and Shang 2007).

Our question of interest is in the same class as the inaccurate inventory problem. As in that problem, the correct inventory level is not known, but the cause is a lack of information rather than erroneous information. The end result is very similar, in that elimination of the cause provides for correct information regarding inventory level, which leads to substantial profitability increases.

There is a separate but related literature on the VOI. As a general topic, VOI has been explored since the earliest operations research textbooks (e.g., Wagner 1969, ch.16). Furthermore, there has been recent interest in this topic by both practitioners and academics.

The inventory-based VOI literature has come to very different conclusions on the actual numerical value of information. Some show that VOI can be substantial. Gavirneni, Kapuscinski and Tayur (1999) show potential values ranging from 10% to 90%. They studied a supplier-retailer dyad and compared the case in which the supplier forecasts retail orders from past data to the case in which the supplier has full information about the retailer's inventory position. Although all parameters have an effect, the main cause of the wide disparity in results is attributed to supplier capacity and the nature of demand. The VOI was small with tight supplier capacity situations as the managerial decision was the same with or without information (produce constantly). Conversely, when capacity was not restrictive information allowed production to be more selective and therefore the VOI is higher. The nature of the demand distribution also had enormous effects: Changing demand distributions from Erlang to Uniform at high capacity altered VOI from less than 30% to 90%. Cachon and Fisher (2000) study a seemingly similar environment, but with 1 supplier and *N* retailers rather than a dyad. The focus of Cachon and Fisher is more on using information to properly allocate inventory among retailers. In contrast to the large VOI seen in Gavirneni, et al. (1999), Cachon and Fisher find an average 2% gain from the supplier having knowledge of retailer inventory positions.

In keeping with the two studies above, parameterization and the focus of inquiry account for large differences in other VOI studies. As examples of the order of magnitude of different findings, Simchi-Levi and Zhao (2003, 2004) show cost reductions of 5%-35%, Moinzadeh (2002) demonstrates that different parameterization leads to VOI with a range from 0% to 23%. Gaur, Giloni and Seshadri (2005) report that sharing information leads to an average reduction in manufacturer's safety stock of 16%, but that the benefit is highly dependent on parameterization. In a meta-analysis of the VOI in inventory literature, Ketzenberg, Rosenzweig, Marucheck and Metters (2007) identify 27 academic manuscripts that numerically calculate the VOI in supply chain inventory settings. Their findings are that VOI is not a simple concept and that results can be idiosyncratic. For example, "Of the twelve studies that report sensitivity with respect to the CV (of demand), four report that VOI decreases with respect to increases in the CV, six report that VOI increases with respect to increases in the CV, and the remaining two report that VOI is largest at intermediate values of the CV" (Ketzenberg, et al. 2007, p. 1235). Ketzenberg, et al.

(2007) summarized the literature by stating that VOI tends to be larger when there is greater uncertainty or sensitivity to uncertainty due to product obsolescence or perishability. The authors note that structural choices, such as periodic review (as opposed to continuous review) can heighten VOI, and larger numbers of facilities in the system (e.g., supplying *N* retailers versus one) can reduce VOI due to risk pooling effects. Hence, realistic parameterization is essential in VOI problems.

We use the typical method to measure VOI: we compare a "base scenario" with a given set of information to an "information scenario" that is structurally identical to the base scenario but with additional information available. VOI is defined as the percentage improvement that a system achieves through the use of additional information relative to a base scenario. As shown from the prior manuscripts, accurate parameterization is essential to deriving realistic VOI. Parameters used in our experiments are from real data obtained from several independent sources.

Due to the inventory policies used, this work is also related to literature on so-called (*R,S*) or (*R,Q*) policies, as noted previously. Inventory work in this area typically attempts to solve for optimal order-up-to or order quantity amounts simultaneously with the optimal replenishment period (e.g., Zheng and Federgruen 1991). Given the fixed order-up-to amount of filling the machine, such complexities are not needed in this study. Further, inventory theorists working on this issue focus on the other side of the problem from what is required here. In the words of Silver, Pyke and Petersen (1998, p.288), "The value of *R* is assumed to be predetermined" by convenience, not design, and the difficulty is in calculating the order-up-to points. Here, the order-up-to point is a forced amount, and calculating *R* is of interest.

Maxwell and Muckstadt (1985) and Federgruen and Zheng (1993) considered *R* to be a decision variable used to coordinate a multi-echelon production-distribution network for a single product. Each study assumes constant demand and develops power-of-two policies. Similarly, Roundy (1986) coordinates a multi-echelon, multi-product production / inventory system with constant demand by using power-of-two policies. Evans (1967) studied inventory management for multiple products in a capacitated production environment. He presented a base-stock policy (also known as an *S* policy) that allows for stochastic demand and treats the length of the replenishment interval as a given parameter. Our work follows directly from these by seeking the optimal replenishment interval R^*.

Literature on the inventory aspects of vending machines is not large and does not address the issue at hand. You (2005) models the downstream purchasing decision for owners of large numbers of vending machines in the face of quantity discounts. Maitra and Dalal (2001) investigate how frequently individual machines should be visited in the absence of information. Anupindi, Dada, and Gupta (1998) look at product substitution in cold beverage vending machines, but do not consider the inventory stocking decision.

4 Single Product Model

Our ultimate goal is to model VOI for a realistic situation. However, modeling many potential products per machine and thousands of machines in hundreds of locations per vending firm is taxing on current computing. Consequently, we first simplify the problem for the purpose of determining insights, then add more realistic elements in the next two sections.

In this section we examine the VOI that can be garnered by attending to a machine in isolation and treating a machine as though it has only one product. This abstraction is useful in gaining understanding, and we will find that it is an acceptable approximation to more complex environments. In terms of typical inventory issues, treating multiple products as a single product can be seen as representing a lost sales substitution rate of 100%. The opposing substitution extreme of 0% will be examined in the next section (for a recent review of product substitution literature, see Gaur and Honhon 2006. Experimentally based measurements of substitution are found in Ketzenberg, et al. 2000).

The standard operating procedure for most of the unattended point of sale industry is to visit and refill a given machine every R time periods. We call this *static scheduling*. The use of "smart machines" that report their inventory levels remotely allows the servicing firm to visit each machine only when its inventory level requires. This operational mode is called *dynamic scheduling* in the trade. The value of information is measured as the percentage improvement in profit achieved through dynamic scheduling, relative to static scheduling. We first analytically define and examine the static and dynamic scheduling processes, then conduct a numerical study to assess the VOI and perform a sensitivity analysis. For mathematical tractability, we use the Poisson distribution in our analytic work. The numerical study uses discrete distributions appropriate to the parameters.

The base model assumes a single product, i.i.d demand in each period, lost sales, and a fixed replenishment cost. We summarize our main notation in Table 1.

4.1 Static Schedule

We analytically derive the optimal replenishment interval R for a single-product system with stochastic demand and given order-up-to levels. Here each visit includes both inventory review (ascertaining current inventory level) and replenishment (filling the machine). Hence, this is a base stock policy in which $S = Y$.

In general, the function for the system's expected profit per period is

$$\pi_s(R) = m\lambda - \frac{K + L_s(R)}{R},$$

with the expected loss due to stockouts being

$$L_s(R) = (m+p)\sum_{k=Y+1}^{\infty}(k-Y)\phi^R(k) .$$

Specifically, for Poisson demand with mean rate λ per period, we have $\phi^R(k) = \dfrac{e^{-\lambda R}(\lambda R)^k}{k!}$ and

$\Phi^R(x) = \sum_{k=0}^{x}\dfrac{e^{-\lambda R}(\lambda R)^k}{k!}$. Then the profit function reads

$$\pi_s(R) = \lambda m - \frac{K}{R} - \frac{(m+p)}{R}\sum_{k=Y+1}^{\infty}(k-Y)\frac{e^{-\lambda R}(\lambda R)^k}{k!}.$$

Since

$$\begin{aligned}\sum_{k=Y+1}^{\infty}(k-Y)\frac{e^{-\lambda R}(\lambda R)^k}{k!} &= \sum_{k=Y+1}^{\infty}k\frac{e^{-\lambda R}(\lambda R)^k}{k!} - Y\left(1-\Phi^R(Y)\right)\\ &= \lambda R - \lambda R\sum_{k=1}^{Y}\frac{e^{-\lambda R}(\lambda R)^{k-1}}{k-1!} - Y\left(1-\Phi^R(Y)\right)\\ &= \lambda R\left(1-\Phi^R(Y-1)\right) - Y\left(1-\Phi^R(Y)\right)\end{aligned},$$

$$\pi_s(R) = \lambda m - \frac{K}{R} - \frac{(m+p)}{R}\left[\lambda R\left(1-\Phi^R(Y-1)\right) - Y\left(1-\Phi^R(Y)\right)\right]. \quad (1)$$

Note that technically R is an integer, but here we assume a continuous approximation. The derivative of $\Phi^R(x)$ with respect to R is

$$\begin{aligned}\frac{d\Phi^R(x)}{dR} &= \sum_{k=0}^{x}\left[\frac{-\lambda e^{-\lambda R}(\lambda R)^k}{k!} + \frac{e^{-\lambda R}\lambda k(\lambda R)^{k-1}}{k!}\right]\\ &= -\lambda\sum_{k=0}^{x}\frac{e^{-\lambda R}(\lambda R)^k}{k!} + \lambda\sum_{k=0}^{x-1}\frac{e^{-\lambda R}(\lambda R)^k}{k!}\\ &= \frac{-\lambda e^{-\lambda R}(\lambda R)^x}{x!} = -\lambda\phi^R(x).\end{aligned}$$

Therefore,

$$\frac{d\pi_s(R)}{dR} = \frac{K}{R^2} - (m+p)\left(\lambda\phi^R(Y-1) - \frac{RY\phi^R(Y) - Y\left[1-\Phi^R(Y)\right]}{R^2}\right)$$
$$= \frac{K}{R^2} - (m+p)\left(\lambda\phi^R(Y-1) - \frac{\lambda R^2\phi^R(Y-1) - Y\left[1-\Phi^R(Y)\right]}{R^2}\right)$$
$$= \frac{K}{R^2} - (m+p)\left(\frac{Y\left[1-\Phi^R(Y)\right]}{R^2}\right)$$
$$= \frac{K-(m+p)Y\left[1-\Phi^R(Y)\right]}{R^2}.$$

It follows that

$$\frac{d\pi_s(R)}{dR} = 0 \qquad \Leftrightarrow \qquad \Phi^R(Y) = 1 - \frac{K}{(m+p)Y}.$$

That $d\Phi^R(Y)/dR < 0, \forall R,$ implies that there is at most one solution R^* to $\Phi^R(Y) = 1 - K/(m+p)Y$. That $\Phi^R(Y) \to 1$ as $R \to 0$, and $\Phi^R(Y) \to 0$ as $R \to \infty$ implies that such an R^* exists if $K \le (m+p)Y$, which is obviously necessary for this to be a profitable business model. If $d\pi/dR = 0$, then $d^2\pi/dR^2 < 0$, so this point is a maximum. Therefore, there exists a unique R^* that maximizes profit, and this R^* is the lone critical point of $\pi_s(R)$. This simplifies the computational study's search for the optimal replenishment period (which must be an integer): it only needs to consider the two integers $\lfloor R^* \rfloor$ and $\lceil R^* \rceil$. Additionally, note that R^* is increasing in fixed replenishment cost K and decreasing in $(m+p)$ because increasing $(m+p)$ makes stock-outs more expensive.

Substituting for $\Phi^R(Y)$ in equation (1) yields

$$\pi_s(R^*) = \lambda(m+p)\Phi^{R^*}(Y-1) - p\lambda.$$

Since $\Phi^R(Y) = \Phi^R(Y-1) + \phi^R(Y)$, we have $\Phi^{R^*}(Y-1) = 1 - \frac{K}{(m+p)Y} - \phi^{R^*}(Y)$ and optimal profit

$$\pi_s(R^*) = \lambda m - \frac{\lambda K}{Y} - \lambda(m+p)\phi^{R^*}(Y).$$

Note that $\phi^R(Y)$ is strictly positive in *Y*, so $\lambda m - \frac{\lambda K}{Y}$ is an upper bound for all *Y*>0 and, since $\phi^R(Y) \to 0$ as $Y \to \infty$, the optimal profit approaches the bound as $Y \to \infty$. The upper bound can be interpreted as the maximum profit per period in a completely deterministic system in which λ units are sold each time period and a replenishment is made at the exact moment that the last item is sold: the replenishment period would be Y/λ, so the allocation per period of the replenishment cost is $\lambda K/Y$.

4.2 Dynamic Schedule

Now consider the case in which the manager can remotely determine the machine's inventory level at the beginning of each time period and then decide whether or not to replenish it. This is a periodic (*s*,*S*) system. These have been shown to be optimal under a variety of conditions (e.g., Iglehart 1963, Veinott and Wagner 1965, Zheng 1991, Beyer and Sethi 1999). Our implementation is different in that *S* is given and therefore our task now is to determine the optimal value of *s*, say s^*.

The manager's objective is to maximize average profit per period. Following Veinott and Wagner (1965), but adjusting for lost sales, we define the expected lost sales and penalty costs incurred during a time period that begins with inventory $y \geq 0$ to be

$$L_d(y) = (m+p)\sum_{k=y+1}^{\infty}(k-y)\phi(k),$$

with $\phi(k)$ being the probability that period demand equals *k*. Let *Y* be the starting inventory and $T(D)$, a random variable, be the first period in which the cumulative demand exceeds *D*. Then

$$L(Y,D) = L_d(Y) + \sum_{i=1}^{\infty}\sum_{x=0}^{D} L_d(Y-x)\phi^i(x) \qquad (3)$$

is the expected loss incurred during $T(D)$. Here, $L_d(Y-k)$ is the expected loss in the (i+1)th period since replenishment, given that after i >0 periods cumulative demand equals $k \le D$ with probability $\phi^i(k)$. $L(Y,D)$ is simply the convoluted effect of all possible realizations of D during $T(D)$.

The cumulative probability that over i periods demand does *not* exceed k equals $\Phi^i(k) = \sum_{t=0}^{k}\phi^i(t)$. The renewal function $M(D) = \sum_{i=1}^{\infty}\Phi^i(D)$ then gives the expected number of periods before cumulative demand exceeds D. Choosing $D = Y - s$ in (3), $L(Y, Y-s)$ denotes the expected cycle costs between replenishments, and $T(Y-s)$ denotes the cycle time with expectation $M(D)$+1. The average profit per period equals $\pi_d(s) = m\lambda - \dfrac{L(Y,Y-s)+K}{1+M(Y-s)}$, so s^* maximizes this expression.

Because we cannot find a closed-form expression for s*, we may use the following bounds to limit the search space (adapted from Veinott and Wagner 1965). Let $\underline{s}$ be the smallest integer for which $L_d(\underline{s}) \le L_d(Y) + K$, then $\underline{s} \le s^* \le Y-1$. Additionally, we approximate the optimum profit and present bounds on that approximation. Suppose that the starting inventory in the last period of a cycle equals exactly s. The lost sales and penalty costs in a cycle are incurred in the last period, so $L(Y,Y-s) \approx L_d(s)$, and the cycle time is approximately $1+(Y-s)/\lambda$. Hence, an approximation for average profit per period is

$$\pi_d(s) \approx \tilde{\pi}_d(s) = m\lambda - \frac{L_d(s)+K}{1+(Y-s)/\lambda}, \qquad (4)$$

which is readily calculated. If s^* is optimal, then naturally $\tilde{\pi}_d(s^*) \ge \tilde{\pi}_d(s^*-1)$ and $\tilde{\pi}_d(s^*) \ge \tilde{\pi}_d(s^*+1)$. Using standard mathematics this can be rewritten as

$$\lambda(m+p)\Phi(s^*-1) - \lambda p \le \tilde{\pi}_d(s^*) \le \lambda(m+p)\Phi(s^*) - \lambda p. \quad (5)$$

Equation (5) holds for $s^* \ge 0$ with $\Phi(-1) = 0$.

Recall that when period demand is Poisson, the static case has maximum profit

$$\tilde{\pi}_s(R^*) = \lambda(m+p)\Phi^{R^*}(Y-1) - \lambda p, \quad (6)$$

so combining (5) and (6) leads to the following bounds on the approximate value of information:

$$\lambda(m+p)\left(\Phi(s^*-1) - \Phi^{R^*}(Y-1)\right) \le \tilde{\pi}_d(s^*) - \pi_s(R^*) \le \lambda(m+p)\left(\Phi(s^*) - \Phi^{R^*}(Y-1)\right). \quad (7)$$

While the bounds for the VOI in equation (7) provide closed form expressions to assess model sensitivity, they are based on the difference $\Phi(s^*) - \Phi^{R^*}(Y-1)$, which complicates the analysis. Even so, by examining how s^* and R^* change with respect to the parameters we can then also assess how the difference $\Phi(s^*) - \Phi^{R^*}(Y-1)$ changes. A decrease (increase) in the difference corresponds to a decrease (increase) on approximate VOI. Unfortunately, a close evaluation of the expressions enables us to only assess the sensitivity with respect to two parameters: the VOI increases with respect to λ and decreases with respect to Y. For the other parameters (e.g. K, m, p) a change in their value will result in both $\Phi(s^*)$ and $\Phi^{R^*}(Y-1)$ moving in the same direction so that the net difference $\Phi(s^*) - \Phi^{R^*}(Y-1)$ cannot be unequivocally ascertained. Hence, we clarify our understanding through a numerical study.

4.3 Numerical Study

We conduct a numerical study to evaluate both the magnitude and sensitivity of the VOI across a broad range of practical operating environments. This study is predicated on a factorial design of the parameter values summarized in Table 2. For each permutation of the parameter values, we compute the optimal expected profit for each information case using the respective optimal policies. Parameter values are based on data compiled from a variety of sources. We obtained two years' data for 55 snack machines that constitute a typical route for a vending machine operator in Texas. The data include sales since the last visit, date of visit, product cost, product price, capacity, and units restocked, for each item in each machine for each visit. We physically accompanied drivers on routes to obtain direct data on transportation times between and within customers' sites, machine stocking times, and general industry practices. Our data were corroborated with national data from the trade association (NAMA 2008) and other vending operators that we interviewed.

Demand $\phi(\cdot)$ has mean periodic demand values of 6, 12, and 24. These values correspond to the 10th, 50th, and 90th percentile of our data. The 50th percentile was corroborated with national data. Demand uncertainty is captured by the coefficient of variation (σ / μ) with values of 1.0, 1.5, 2.0, and 2.5, which represent the full spectrum of our data. Given the nature of the problem, the demand distributions used should be discrete and not allow negative values. Consistent with the definitional boundaries of the distributions, we use truncated binomial, Poisson, and negative binomial distributions corresponding to different variance/mean ratios. See Agrawal and Smith (1996) regarding the realism of assuming negative binomial distributions for demand in retail settings. The profit margins ranging from \$0.4 to \$0.7 represent the 20th to 80th percentile of empirical data of 167 products typically for sale. Setup costs of \$5 through \$8 are based on average times for servicing and wage rates obtained from surveys of the industry by the national association (NAMA 2008), which was corroborated by our physical observations of route drivers. Capacities correspond to those of commonly used machines. Finally, penalty costs for lost sales are chosen so that the values cover a broad range of service requirements. The penalty cost for an individual lost sale may be near zero, but if many lost sales occur it is believed that the vendor will be replaced; this indicates a long term lost sale cost that may be large.

The VOI represents the percentage difference in expected profit that is realized through the use of information, relative to the static schedule case. Exact numerical solutions are obtained for each information case. In Table 3, we report results at given percentiles for the 1,728 experiments for VOI (profit) along with measures for the improvements in the service level and the reduction in the number of visits. For example, the row labeled 0.05 under percentile contains the results of three different experiments. Under "Profit Increase" is the experiment that had the 5th percentile of VOI, under "Visit Decrease" is the experiment that contained the 5th percentile in visit reduction (all experimental cells resulted in reducing the number of visits), and under "Service Increase" is the experiment that contained the 5th percentile of service level increase.

From Table 3, two observations emerge: 1) information enables a substantial reduction in machine visits and a modest improvement in service (fill rate) and, 2) the range of VOI shows demonstrable sensitivity to model parameterization that depends largely on system behavior, as we discuss below, with the 5th percentile of profit improvement being 1.1% and the 95th percentile of improvement being 10.7%. Service increases are small, with the 95th percentile only achieving a 3.1% increase in service level (fill rate). The small improvements in service are due to the very high service levels sustained by the traditional policy, with 99% fill rates not unusual in the optimal static visit schedule. These high fill rates are not merely a modeling artifact – they were also observed when shadowing route drivers. The high fill rates stem from the relatively low cost of servicing a machine compared to far higher potential lost sales cost. The main benefit from information is knowing when not to visit a machine, with the median visit decrease being 38.0%.

We summarize our sensitivity analysis in Table 4, which provides average values across all experiments with the fixed parameter value identified by the row label. For example, the average VOI (profit increase) is 6.8% when machine capacity is 320. Note that the VOI is largest in cases where the level of profitability is poor in the static case (as profit margins increase to infinity or capacity diminishes to zero, both policies converge to the same solution: restock every time period so there would be no VOI). That is, the VOI is largest for high cost environments with limited capacity machines that are subject to high demand uncertainty. The most volatile parameter for achieving profit increases is the CV of demand. A CV of 1.0 averages only a 1.7% increase in profit, while a CV of 2.5 averages a 9.6% profit increase. The sensitivity of the VOI to model parameters is similar to other studies in this area, with one difference. Earlier works (e.g. Gavirneni et al. 1999) show that the VOI decreases as capacity decreases. However, those studies consider manufacturing capacity. They show that with lower levels of capacity, a supplier is unable to react to information since it would already be producing continuously, i.e., at capacity. In the current study it is inventory that is capacitated. As capacity decreases, inventory and, hence, sales are both increasingly constrained. It follows that with information, there are increasing opportunities to reduce lost sales.

5 Multiple Product Model

In the last section we modeled the vending machine as though it had just one product. This would be appropriate in practice if substitutability among products was very high. Here we view the other extreme by assuming that there is no substitutability between machine slots. It is common in drink machines, but not as common in snack machines, to have multiple slots dedicated to a single product. Consequently, a machine with partial substitutability would be the most accurate way to model practice. However, by bracketing the reality of partial substitutability with the two extremes of no substitutability and complete substitutability, we bound the results that would be obtained in reality. As will be seen later, the VOI is highly similar between these two scenarios.

Below, we first introduce revised policies for the case of multiple products and then perform a sensitivity analysis through a numerical study based on industry data.

5.1 Static Schedule

For n products, the model is perfectly analogous to the one-product case: the function for the system's expected profit is

$$\pi(R) = -\frac{K}{R} + \sum_{i=1}^{n}\left[\lambda_i (m_i + p_i)\Phi_i^R(Y_i - 1) + \frac{Y_i}{R}(m_i + p_i)\left[1 - \Phi_i^R(Y_i)\right] - p_i\lambda_i\right],$$

where each product i has Poisson demand with a mean rate per period of λ_i and a stocking level of Y_i. Let m_i be product i's profit margin and p_i be product i's penalty cost. Therefore,

$$\frac{d\pi(R)}{dR} = \frac{K}{R^2} - \frac{1}{R^2}\sum_{i=1}^{n} Y_i(m_i + p_i)\left[1 - \Phi_i^R(Y_i)\right]$$

Setting this equal to zero yields

$$\sum_{i=1}^{n} Y_i(m_i + p_i)\left[1 - \Phi_i^R(Y_i)\right] = K. \qquad (8)$$

Furthermore,

$$\begin{aligned}\frac{d^2\pi(R)}{dR^2} &= -\frac{2K}{R^3} + \frac{2}{R^3}\sum_{i=1}^{n} Y_i(m_i + p_i)\left[1 - \Phi_i^R(Y_i)\right] - \frac{1}{R^2}\sum_{i=1}^{n} Y_i(m_i + p_i)\lambda_i\phi_i^R(Y_i) \\ &= -\frac{1}{R^2}\sum_{i=1}^{n} Y_i(m_i + p_i)\lambda_i\phi_i^R(Y_i) < 0,\end{aligned}$$

if $d\pi(R)/dR = 0$. Therefore, the R that solves (8), which can easily be found numerically, maximizes expected profit.

5.2 Dynamic Schedule

Just as in the single product case, dynamic scheduling in the multi-product case is predicated on the decision of whether to visit the machine in the current period or to delay the decision until the next period. Unfortunately, the inventory policy is no longer of the (s, S) variety, and the problem of deriving an optimal policy becomes analytically intractable. Instead, we develop a heuristic policy. This heuristic, while not optimal, performs extremely well, as we later show.

The structure of the heuristic is based on the tradeoff between the lost sales that arise from making a visit too late and the cost of making a visit too early. Specifically, the decision to visit in a period occurs when the marginal benefit of making a visit is positive. Let I_i denote the beginning inventory level for product i in the current period. Then the benefit of visiting is the expected reduction in lost sales that would otherwise occur without a visit. Here, the benefit is given by

$$\sum_{i=1}^{n}(m_i + p_i)\left[\sum_{x=I_i+1}^{Y_i}(x - I_i)\phi(x)\right].$$

The marginal cost of making a visit in one period versus in some future period is predicated on the allocation of the cost of making a visit (K) over an unknown expected replenishment period. This is similar in structure to the single product, dynamic case in which the expected replenishment period is given by $1+(Y-s)/\lambda$ (see equation (4)). However, since the expected replenishment period is unknown in the multi-product case, we use the optimal replenishment period R^* of the static schedule case as a proxy. Hence, the decision to make a visit in the current period will occur if

$$\sum_{i=1}^{n}(m_i + p_i)\left[\sum_{x=I_i+1}^{Y_i}(x - I_i)\phi(x)\right] - \frac{K}{R^*} > 0 .$$

Intuitively, the VOI is determined relative to the static case, so the cost allocation of K over R^* periods is a logical extension. We tested the heuristic in the one product case for all 1,728 experiments. Across *all* experiments, there were no differences in the VOI from the results reported in Section 4 that were obtained with the optimal policy. In the next section we show that the VOI calculated by using the heuristic for a machine with multiple products is very similar to that observed for the single product case.

5.3 Numerical Study

A typical snack vending machine holds approximately 40 different products. We obtained empirical demand data at the individual product level on 55 machines over a period of nearly two years, as described earlier. Here, we model three total machine demand levels of 6, 12, and 24, representing demand from the 10th, 50th, and 90th percentiles of our data, just as was done in section 4. Demand means for individual products are shown in Table 5. These were obtained by averaging the parameter values of the products of several similar machines, then using smoothing to obtain the values shown.

Empirically, on the individual product level, the coefficient of variation was noted to escalate inversely to the demand mean. This was modeled in our experiments by assigning the coefficients of variation in Table 6 (the same as used in the experiments in section 4) to the highest demand product in the machine. The variance to mean ratio of that product was then used to assign variance parameters to all other products. As an example, for the total demand = 12, CV = 1.5 scenario, the highest demand product has a mean demand of 0.8 per period. A CV of 1.5 dictates a standard deviation of 0.8 x 1.5 = 1.2. This translates into a variance/mean ratio of (1.2 x 1.2)/0.8=1.8. The variance/mean ratio is kept constant for all other products. The 20^{th} highest demand product has mean demand of 0.3 per period, so it is assigned a variance of [variance to mean ratio 1.8 multiplied by the mean] 1.8 x 0.3 = 0.54.

An additional parameter in a multi-product situation that does not occur in a single-product situation is the correlation of demand between products. Empirically, demand correlation was found to be highly positive. In our data, the average correlation of demand from individual products with overall machine demand ranged from nearly 0.6 to nearly 0.8. Consequently, we use correlations of 0.6 and 0.8 to cover the extremes of practice.

Reasons for high demand correlation are speculative, but not difficult to imagine. For a machine in a public arena with an event cancelled, the demand for all products will be zero. If the public facility is populated with twice the normal population for a highly popular event, then demand is likely to be higher for all products. In a factory setting, an early shift shut down or even a birthday celebration that causes large numbers of employees to leave the building for lunch causes all products to have less demand. A factory requiring overtime or a double shift will increase demand. Vending companies are not generally notified of such events.

Tests were run over values similar to those of the single product case for profit margin, lost sales penalty cost, machine setup cost, and machine capacity (Table 2), resulting in 3,456 experiments. Note that the capacities specified per product correspond to machine capacities of 320, 400, and 480, which correspond to the capacities in the prior experiments that treat the machine as having a single product.

Unfortunately, exact analysis becomes computationally infeasible with multiple products. Here, we employ a simulation study to generate results for analysis. We simulated the inventory performance for both static and dynamic scheduling. Each of the 3,456 experiments is simulated for 2,100 periods and replicated 30 times. The first 100 periods of each replication are set aside as the simulation warm-up period so that statistics are calculated for 2,000 periods in each replication. The warm-up period was chosen for convenience, and to be in keeping with the simulation research tradition. (As the system resets for each replenishment, a warm-up period is technically not necessary.) In each replication, the random number streams across all experiments are identical in order to reduce the sampling error. The standard deviations of the mean results obtained via simulation are sufficiently small that the statistics for our results are omitted.

Summary results by percentile of simulated cases are presented in Table 7. The results are similar to those obtained in the single product case, but the single product case (Table 3) consistently understates the VOI compared to modeling a multi-product environment. For the 50^{th} percentile of results (median), treating the machine as though composed of only one product brought a 3.4% profit increase, while modeling all 40 products separately yields a 4.5% profit increase. The reasons behind profit increases are also similar: a small increase in service level (fill rate) and a large decrease in the number of machine setups (visits).

Sensitivity of the VOI to model parameterization (Table 8) is also very similar to that observed for the single product case (Table 4). The VOI is increasing in the aggregate level of machine demand. The effect of increased demand is the same as decreasing capacity – the more tightly capacitated the system, the greater the opportunity to improve service through information - which is the same result observed for the single product case. The VOI is highest in high operating cost environments (low margins, high penalty cost, and tight capacity). Conversely, the lowest VOI environments are those with no penalty cost for lost sales and large capacity machines. The combination of those two parameter values reduces the influence of lost sales and fewer replenishments, so the opportunity to both increase sales and reduce replenishments through information is lessened.

As can be seen by comparing Tables 4 and 8, modeling the machine as a single product has very similar, but consistently more modest, results in every parameter setting as modeling the machine with 40 products. Note that the correspondence arises even though the single product case assumes an aggregate analysis of a machine in which all products are 100% substitutable, whereas the multi-product case assumes zero percent product substitution. As such, we are confident that the additional complexities associated with modeling product substitution are unnecessary from a managerial perspective and that the simpler aggregate analysis is sufficient for decision making. Because of this, we will use the simpler, single product model in extending our evaluation of the VOI to the more complex, multiple machine case.

6 Multiple Machine Model

In a practical setting, machines are not isolated units – they are located in groups. A large customer may have several hundred machines in a corporate campus. Consequently, there are two distinct setups to consider: the transportation time to/from the group of machines (*group setup*) and the individual machine setup explored earlier.

A combination of a lack of information on individual machine/product demand and a desire to rationalize setups has motivated an estimated 90% of the industry to service all machines in a group at the same time every R time periods (NAMA 2008, Maras 2007); i.e., to use static scheduling. The ability to use inventory information via dynamic scheduling is facilitated by an operational change: the ability to visit only those machines in a group that require servicing. Here, we compare the static scheduling (no information) environment to the dynamic scheduling environment to determine VOI.

It could be argued that this benefit could be achieved without information. That is, when servicing a group of machines every R time periods, only service those that appear to require it upon the visual inspection and judgment of the operator. However, the practicalities of machine visit scheduling argue against it. Machines in a "group" are still physically separated, often on different floors of a building or in different buildings on a corporate campus. Consequently, some time must be spent on a review, even if replenishment is not made. Further, given the static schedule of visiting a group every R time periods, a machine not replenished with the others of the group must wait an additional R time periods to be replenished, whereas with information, it may be replenished in the next time period. Given the high level of uncertainty in product demand, waiting for the next replenishment cycle risks the cost of substantial stock-outs.

Attending to groups of machines rather than individual machines is reminiscent of the "inventory-routing" literature (e.g., Gaur and Fisher 2004, Metters 1996). Consequently, it might be thought that by engaging multiple machines, a routing algorithm must also play a prominent role. However, there are stark differences in the operating environments that allow us to ignore the routing component for optimization. In the inventory-routing literature, customers have one demand point (e.g., an oil tank to fill with hundreds of gallons), and customers are geographically distant from one another, requiring an operator to drive dozens of kilometers between demand points. For example, the average number of stops per route is usually less than two in Gaur and Fisher (2004). In contrast, in the vending machine environment, 50 machines are routinely serviced in a shift. This is because the group setup physically places the operator at the customer campus, and each campus has multiple demand points that are geographically close enough for the operator to walk between them. Further, the different customer campuses tend to be close enough together—often in the same industrial park–that several are visited in one route.

6.1 Model

Additional notation is required to describe the group versus individual machine setups. It is also slightly altered to reflect that we now have multiple machines, but not multiple products. Let z denote the number of machines and j denote a specific machine $j \in \{1,2,\ldots,z\}$. There are now two setup costs. Let K_G denote the group setup cost, which corresponds to the travel time to a group of machines. Let K_M denote the individual machine setup cost, which corresponds to the time to service a machine. For the static case, the objective function is now expressed as

$$\pi(R) = -\frac{K_G + nK_M}{R} + \sum_{j=1}^{z}\left[\lambda_j\left(m_j + p_j\right)\Phi_j^R\left(Y_j - 1\right) + \frac{Y_j}{R}\left(m_j + p_j\right)\left[1 - \Phi_j^R\left(Y_j\right)\right] - p_j\lambda_j\right],$$

which is structurally identical to that for the multi-product case. As shown in Section 5 for that case, the R that maximizes expected profit can be easily found numerically. We omit the details here.

The dynamic schedule replenishment policy for multiple machines is similar in form to the multi-product heuristic. Here, there are two types of decisions that must be made in each time period: First, determine if there are one or more machines that justify a visit on the basis of the machine setup cost. This justification for a machine j is based on whether

$$\left(m_j + p_j\right)\sum_{x=I_j+1}^{Y_j}\left[x - I_j\right]\phi_j(x) - \frac{K_M}{R_j^*} > 0,$$

where R_j^* denotes the optimal replenishment period for a static schedule visiting machine j in isolation (as per equation (1), where $K = K_M$). Then let V denote the set of machines in which a visit is justified. The next step is to determine if visiting the subset of machines V is warranted, given the group setup cost. We will then visit and replenish these machines if

$$-\frac{K_G}{R^*}+\sum_{j\in V}\left[\left(m_j+p_j\right)\sum_{x=I_j+1}^{Y_j}\left[x-I_j\right]\phi_j\left(x\right)-\frac{K_M}{R_j^*}\right]>0,$$

where R^* denotes the optimal replenishment period for the multi-machine static case.

6.2 Numerical Study

Just as with the multi-product case, we employ a numerical study to evaluate the differences between dynamic and static scheduling using simulation. The simulation procedures are also the same as those for the multi-product case and are not repeated here. As in the single machine case, we take an aggregate view of demand for each machine as though for a single product. Parameters are set in accordance with Table 9. The parameters are similar to those used in the prior studies and provide 162 experimental cells for each machine group size. Because low demand variance and high levels of capacity provided very little VOI in the prior experiments, we focus on the higher levels of CV and tighter capacities in this set of experiments.

We investigate machines in groups of 1, 5, 10, 20 and 40. While a one-machine "group" is not realistic, it is valuable as a benchmark for clarifying certain results.

Summary results are reported in Table 10. The median VOI is 9.2% for the one-machine scenario, declining to 6.1% for a 10 machine group, then increasing for 20- and 40-machine groups. The convexity in the number of machines per group is due to the two opposing setup rationalizations. VOI is highest in the one–machine-per-group case because both setups can be avoided when the machine is still relatively full. In the 5 machine scenario the group setup is made more frequently in the full information case, while fewer individual machine setups are made. In the 40 machine scenario, the group setup is made in virtually every time period in the full information case, but substantial savings are made by having fewer individual machine setups. The competing setup rationalizations mean that with a group of only 5 machines, the VOI is smaller, though still significant.

For completeness, we report a sensitivity analysis in Table 11. The results are similar with those reported for the previous studies (Tables 4 and 8). As before, the highest VOI is attained when demand is more uncertain and the lost sales penalty cost is high. Note that the VOI is decreasing slightly with respect to the group setup cost (excluding the one machine case) and increasing with respect to the machine setup cost. As the group setup cost increases, the system becomes more constrained in its ability to take advantage of information provided by individual machines. At the limit, all machines must justify a visit, which approaches the replenishment periodicity of the static case. The increase of the VOI with respect to the machine setup cost corresponds to the increased savings from not visiting individual machines that do not require service under the decision rule for a dynamic schedule.

7 Conclusion

In relation to the VOI literature, unattended points of sale represent an extreme information environment. In the typical supply chain VOI environment, supply chain partners may be distrustful of one another and reluctant to share information, but the entities studied are sentient and are presumed to make rational, if perhaps selfish, decisions. Unattended points of sale, by nature, cannot make any decisions, rational or otherwise. Traditionally, information could not be passed along. However, new technologies allow information to be shared.

Our experiments indicate that monitoring inventory remotely, exemplified by vending machines, is complex. Traditional economic theory would posit that more information would clearly lead to better results. We show that under certain business environments, those characterized by low demand, low demand volatility, or high profit margins, additional information on inventory level has a low impact. However, under conditions of high demand, high demand volatility, and more severe penalties for lost sales, there is potential to significantly increase profitability by making use of information. It should be noted that our estimates are conservative, as they do not include the known benefits associated with reduced employee theft and reduced service time per machine.

However, to be effective, the value of this information must be coincident with business practice changes. The static scheduling environment of visiting groups of machines every R time periods is an ingrained practice in the industry. To capture the potential VOI this business practice must change.

Acknowledgements

We would like to thank Bill Stein for his help with the analytical development.

References

Agrawal, N., S. Smith. 1996. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Research Logistics*, 43(6), 839-861.

Anupindi, R., M. Dada, S. Gupta. 1998. Estimation of consumer demand with stock-out based substitution: An application to vending machine products, *Marketing Science,* 17(4), 406-423.

Balakrishnan, A., M. Pangburn, E. Stavrulaki. 2004. "Stack them high, let 'em fly": lot-sizing policies when inventories stimulate demand. *Management Science*, 50(5), 630-644.

Bell, W., L. Dalberto, M. Fisher, A Greenfield, R. Jaikumar, P. Kedia, R. Mack, P. Prutzman. 1983. Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer. *Interfaces*, 13(6), 4-23.

Beltrami, E., L. Bodin, 1973. networks and vehicle routing for municipal waste collection. *Networks*, 4, 65-94.

Beyer, D., S.P. Sethi. 1999. The classical average-cost inventory models of Iglehart and Veinott-Wagner revisited. *Journal of Optimization Theory and Applications*, 101(3), 523-555.

Borin, N., P. Farris, J. Freeland. 1994. A model for determining retail product category assortment and shelf space allocation. *Decision Sciences*, 25(3), 359-384.

Bultez, A., P. Naert. 1988. SH.A.R.P. Shelf allocation for retailers' profit. *Marketing Science*, 7(3), 211-232.

Cachon, G. P., M. Fisher 2000. Supply chain inventory management and the value of shared information. *Management Science* **46**(8), 1032–48.

DeHoratius, N., A. J. Mersereau, L. Schrage. 2008. Retail inventory management when records are inaccurate. *Manufacturing & Service Oper. Management,* **10**(2), 254-277.

DeHoratius, N., A. Raman. 2008. Inventory Record Inaccuracy: An Empirical Analysis. *Management Science*, **54**(4), 627-641.

de Kok A.G., K.H. Van Donselaar, T. Van Woensel. 2008. A break-even analysis of RFID technology for inventory sensitive to shrinkage, *International Journal of Production Economics,*112(2), 521-531.

Dror, M., M. Ball, 1987. Inventory/routing: reduction from an annual to a short-period problem. *Naval Res. Log.*, 34, 891-905.

Evans, R.V. 1967. Inventory control of a multiproduct system with a limited production resource. *Naval Research Logistics*, **14**(2), 173-184.

Federgruen, A., Y. Zheng. 1993. Optimal power-of-two replenishment strategies in capacitated general production / distribution networks. *Management Science*, **39**(6), 710-727.

Gaur, V., M. Fisher, 2004. A periodic inventory routing problem at a supermarket chain, *Operations Research*, **52**(6), 813-822.

Gaur, V., A. Giloni, S. Seshadri. 2005. Information sharing in a supply chain under ARMA demand, *Management Science*, **51**(6), 961-969.

Gaur, V., D. Honhon. 2006. Assortment planning and inventory decisions under a locational choice model, *Management Science*, **52**(10), 1528-1543.

Gavirneni, S., R. Kapuscinski, S. Tayur. 1999. Value of information in capacitated supply chains, *Management Science*, **45**(1), 16–24.

Iglehart, D. 1963. Dynamic programming and stationary analysis of inventory problems. *Multistage Inventory Models and Techniques*, H. Scarf, D. Gilford, M. Shelly, eds., Stanford University Press, Stanford, CA, 1-31.

Ito, M. 2007. Stroll the streets of vending machine heaven. *Japan Times Online*, April 15, http://search.japantimes.co.jp/cgi-bin/fl20070415x2.html, accessed April 14, 2009.

Kang, Y., S. Gershwin. 2005. Information inaccuracy in inventory systems: stock loss and stockout. *IIE Transactions*, 37, 843-859.

Ketzenberg, M., R. Metters, V. Vargas. 2000. Inventory policy for dense retail stores. *Journal of Operations Management*, 18(3), 303-316.

Ketzenberg, M., E. Rosenzweig, A. Marucheck, R. Metters. 2007. Information value for supply chain inventory replenishment. *European Journal of Operational Research*, 182(3), 1230-1250.

Kok, G., K. Shang. 2007. Inspection and replenishment policies for systems with inventory record inaccuracy. *Manufacturing & Service Operations Management,* 9(2), 185-205.

Lee, H., O. Ozer. 2007. Unlocking the value of RFID. *Production and Operations Management*, 16(1), 40-64.

Maitra, R., S. Dalal. 2001. Pay phones, parking meters, vending machines, and optimal Bayesian decisions on collection times. *Journal of the American Statistical Association*. 96(454), 476-487.

Maras, E. 2007. State of the Vending Industry Report. *Automatic Merchandiser*, August, 43-71.

Maras, E. 2010. State of the Vending Industry Report. *Automatic Merchandiser*, June/July, 32-48.

Maxwell, W.L., J.A. Muckstadt. 1985. Establishing consistent and realistic reorder intervals in production-distribution systems. *Operations Research*, **33**(6), 1316-1341.

Metters, R., 1996. Interdependent transportation and production activity at the United States Postal Service. *Journal of the Operational Research Society*, 47(1), 27-37 .

Milgrom, P., J. Roberts. 1988. Communication and inventory as substitutes in organizing production. *Scandanavian Journal of Economics*, **90**(3), 275-289.

Moinzadeh, K. 2002. A multi–echelon inventory system with information exchange. *Management Science*, **48**(3), 414–426.

NAMA 2008. Personal communication from National Automatic Merchandising Association executives.

Raman, A., N. DeHoratius, Z. Ton. 2001. Execution: the missing link in retail operations. *California Management Review*, **43**,(3), 136-152.

Roundy, R. 1986. A 98%-effective lot-sizing rule for a multi-product, multi-stage production / inventory system. *Mathematics of Operations Research*, **11**(4), 699-727

Silver, E., D. Pyke, R. Peterson 1988. *Inventory Management and Production Planning and Scheduling*, 3rd ed., Wiley & Sons, NY.

Simchi–Levi, D., Y. Zhao 2003. The value of information in a two–stage supply chain with production capacity constraints. *Naval Research Logistics*, 50(8), 888-916.

Simchi–Levi, D., Y. Zhao 2004. The value of information in a two–stage supply chain with production capacity constraints: the infinite horizon case. *Probability in the Engineering and Informational Sciences*, 18, 247-274.

Ton, Z., A. Raman. (forthcoming). The Effect of Product Variety and Inventory Levels on Retail Sales: A Longitudinal Study, *Production and Operations Management.*

Veinott, A., H. Wagner. 1965. Computing optimal (s, S) policies. *Management Science*, 11, 525-552.

Wagner, H.M. 1969. *Principles of Operations Research*, Englewood Cliffs, NJ: Prentice-Hall, Inc.

You, P. 2005. A heuristic approach for multiple item and location ordering problem with quantity discount and capacity constraint. *Journal of the Operational Research Society*, 56(3), 307-315.

Zheng, Y. 1991. A simple proof for optimality of (s,S) policies in infinite-horizon inventory systems. *Journal of Applied Probability*, 28(4), 802-810.

Zheng, Y., A. Federgruen. 1991. Finding optimal (s,S) policies is about as simple as evaluating a single policy. *Operations Research*, 39(4), 654-665.

Table 1: Model Notation

R	Replenishment period for static scheduling
K	Fixed replenishment cost per visit
Y	Machine capacity
m	Margin per unit of sales
p	Penalty cost per unit of lost sales
λ	Mean demand rate
$\phi(.)$	Demand density function
$\phi^R(.)$	R-fold convolution of demand
$\Phi^R(.)$	Cumulative of the R-fold convolution of demand

Table 2: Design of experiments (single product case)

Parameter	Values
Mean Demand per period	6, 12, 24
Demand CV	1, 1.5, 2, 2.5
Margin per unit of sales	0.4, 0.5, 0.6, 0.7
Penalty cost per unit of lost sales	0, Margin, 2*Margin
Setup cost	5, 6, 7, 8
Machine Capacity	320, 400, 480

Table 3: Summary Results (single product case)

	% Change with Information		
Percentiles	Profit Increase	Visit Decrease	Service Increase
0.05	1.1%	18.8%	0.0%
0.25	2.0%	30.0%	0.0%
0.50	3.4%	38.0%	0.1%
0.75	6.1%	48.2%	0.4%
0.95	10.7%	62.0%	3.1%

Table 4: Sensitivity analysis (single product case)

Parameter	Value	Profit Increase	Visit Decrease	Service Increase
Capacity	320	6.8%	39.1%	0.9%
	400	4.8%	39.2%	0.5%
	480	3.5%	38.9%	0.4%
Demand	6	1.5%	32.2%	0.3%
	12	2.3%	39.8%	0.4%
	24	4.3%	45.4%	1.2%
CV	1.0	1.7%	24.9%	0.1%
	1.5	3.2%	35.7%	0.1%
	2.0	5.6%	43.6%	0.2%
	2.5	9.6%	52.2%	2.1%
Margin	0.4	6.1%	36.5%	0.8%
	0.5	5.2%	38.6%	0.6%
	0.6	4.6%	40.0%	0.5%
	0.7	4.2%	41.3%	0.5%
Penalty	0	3.9%	34.0%	0.9%
	1	5.2%	40.0%	0.5%
	2	6.0%	43.3%	0.4%
Visit Cost	5	4.2%	41.2%	0.5%
	6	4.8%	39.7%	0.6%
	7	5.3%	38.4%	0.7%
	8	5.9%	37.1%	0.7%

Table 5: Individual product mean demand (multi-product case)

Individual Product	Demand Scenario		
	Low	Medium	High
Highest Demand	0.5	0.8	1.5
5th Highest Demand	0.3	0.5	1.0
20th Highest Demand	0.1	0.3	0.5
40th Highest Demand	0.1	0.1	0.3
Sum of 40 Products	6	12	24

Table 6: Design of experiments (multi-product case)

Parameter	Values
Mean Demand per period	6, 12, 24
Demand CV of highest demand product	1, 1.5, 2, 2.5
Margin per unit of sales	0.4, 0.5, 0.6, 0.7
Penalty cost per unit of lost sales	0, Margin, 2*Margin
Setup cost	5, 6, 7, 8
Capacity per product	8, 10, 12
Demand correlation between products	0.6, 0.8

Table 7: Summary Results (multi-product case)

	% Change with Information		
Percentiles	Profit Increase	Visit Decrease	Service Increase
0.05	0.9%	4.4%	(1.5)%
0.25	2.1%	15.8%	0.1%
0.50	4.4%	27.9%	0.6%
0.75	9.8%	43.4%	1.8%
0.95	28.6%	63.8%	6.1%

Table 8: Sensitivity Analysis (multi-product case)

Parameter	Value	Profit Increase	Visit Decrease	Service Increase
Capacity	8	13.7%	28.9%	1.6%
	10	7.9%	30.0%	1.2%
	12	5.8%	31.9%	0.8%
Demand	6	4.5%	33.3%	(0.3)%
	12	6.2%	24.4%	1.5%
	24	16.7%	33.2%	2.3%
CV	1.0	2.0%	13.9%	0.4%
	1.5	4.0%	23.9%	0.7%
	2.0	9.1%	35.6%	1.5%
	2.5	21.4%	47.7%	2.2%
Margin	0.4	13.1%	24.0%	1.8%
	0.5	9.3%	28.4%	1.3%
	0.6	7.6%	32.5%	0.9%
	0.7	6.5%	36.2%	0.6%
Penalty	0	4.3%	25.6%	1.8%
	1	8.3%	30.6%	1.1%
	2	14.8%	34.6%	0.6%
Visit Cost	5	6.7%	36.0%	0.7%
	6	8.2%	31.5%	1.1%
	7	9.6%	28.0%	1.4%
	8	12.0%	25.6%	1.7%
Correlation	0.6	8.2%	24.9%	1.3%
	0.8	10.1%	35.6%	1.1%

Table 9: Design of experiments (multi-machine case)

Parameter	Values
Number of machines in a group	1, 5, 10, 20, 40
Mean Demand per period	6 (20% of machines), 12 (60% of machines), 24 (20% of machines)
Demand CV	2, 2.5
Margin per unit of sales	0.45, 0.55, 0.65
Penalty cost per unit of lost sales	0, margin, 2*Margin
Group setup cost	6, 7.5, 9
Machine setup cost	4, 5.5, 7
Capacity per product	320

Table 10: Summary results (multi-machine case)

	Number of Machines				
Percentile	**1**	**5**	**10**	**20**	**40**
0.05	5.4%	3.8%	3.0%	3.4%	3.6%
0.25	7.4%	5.5%	4.5%	4.9%	5.2%
0.50	9.2%	7.2%	6.1%	6.5%	6.7%
0.75	11.5%	9.7%	8.3%	8.6%	8.9%
0.95	15.3%	13.5%	12.0%	12.3%	12.7%

Table 11: Sensitivity analysis (multi-machine case)

		Number of Machines				
Parameter	**Value**	**1**	**5**	**10**	**20**	**40**
Demand CV	2	7.8%	6.2%	5.2%	5.6%	5.8%
	2.5	11.7%	9.4%	8.1%	8.5%	8.8%
Margin	0.45	11.2%	8.7%	7.5%	7.9%	8.2%
	0.55	9.6%	7.7%	6.6%	6.9%	7.2%
	0.65	8.5%	7.0%	5.9%	6.2%	6.5%
Penalty	0	7.5%	5.4%	4.5%	4.9%	5.1%
	Margin	10.0%	8.0%	6.9%	7.2%	7.5%
	2*Margin	11.7%	10.1%	8.7%	9.0%	9.3%
Group Setup	6	8.9%	8.0%	7.0%	7.3%	7.6%
	7.5	9.8%	7.8%	6.7%	7.0%	7.3%
	9	10.6%	7.6%	6.4%	6.7%	7.0%
Machine Setup	4	8.8%	6.3%	5.1%	5.4%	5.7%
	5.5	9.7%	7.8%	6.7%	7.0%	7.3%
	7	10.7%	9.3%	8.2%	8.6%	8.9%