RFID-Enabled Management of Highly-Perishable Inventory:
A Markov Decision Process Approach for Grocery Retailers

Chokdee Siawsolit	Gary Gaukler	Sarun Seepun
Drucker School of Management	Drucker School of Management	Claremont Graduate University
chokdee.siawsolit@cgu.edu	gary.gaukler@cgu.edu	sarun.seepun@cgu.edu

Abstract

We address the challenge of managing perishable inventory. One study was conducted to analyze the effects of recapturing unsatisfied demand, and another to estimate improvements in operational metrics through delaying order placements. Our results indicate that significant profit improvements can be achieved under these scenarios, as evidenced by a greater than 30% median increase in profit margin.

1. Introduction

We take on the perspectives of grocery retailers finding themselves in transition phase toward the online market world. Through two tightly linked studies, we explore how such a retailer might aim to improve the management of their perishables inventory. The main drivers of cost for this product category are outdating (when product reaches its expiration date and cannot be sold any longer) and lost sale events (when stock on hand is not available to fill consumer demand).

In an industry dominated by major players, grocery retailers compete for market share by increasing selection, improving product qualities, and engaging fully in cost-cutting measures in order to offer the lowest prices. Perishable food products have emerged to be a main differentiating factor that draws more customers in [1]. However, the costs associated with having inventory expire on hand presents a challenge in maintaining adequate service level while minimizing operating costs. The faster a product perishes, the higher the chance that outdating costs will occur. The problem can be regarded as a classic tug-of-war between underage and overage costs; with the perishable status adding extra strain on the system by means of having short life times.

Equally a concern, food waste in the distribution chain has been estimated at one-third of the total volume of food produced worldwide [2]. The economic impact of costs incurred can resonate both upstream to producers, and down all the way to consumers. Here, we focus primarily on improving profits and minimizing outdating events for highly perishable food products with very low shelf lives such as ripened fruits and vegetables, or defrosted meat/seafood products. Our overarching question of interest here is: how can we better adapt perishable inventory management to the uncertainty of demand?

A few notable trends that have added pressure, or may offer potential solutions, to the issues above are summarized here. First, the push for healthier eating and wellness concerns are at an all-time high [3]. Wholefoods and Sprouts are examples of retailers well-positioned for this trend; with both having seen rapid growth in the recent years [4]. Health-oriented food shoppers demand less additives in their products; paving the way for organic fruits and vegetables to take up increasing shelf space. Without the help of preservatives or chemicals to prevent deterioration, though, healthier foods may come with a tradeoff of lower shelf lives [5].

In addition, more retailers are adopting some form of radio-frequency identification (RFID) tracking systems. With increased supply chain visibility and data-sharing between suppliers and retailers, now is a prime time to explore and exploit the benefits of having access to a multitude of new information previously not available. Likewise, more consumers than ever have adopted the use of smartphones. This had led to the emergence of additional internet retailing, as well omni-channel experiences through flexible shopping and fulfilment platforms. Many tasks can now be simplified by taking advantage of the copious features that come embedded within recent day smart mobile devices.

We position our work within this broad context of technology-assisted perishable products retailing. The first study involves using a mobile phone application to assist in recapturing unsatisfied demand from a stock out event. Then, we shift our focus to improving inventory performance through delaying the order placement process. Both analyses are performed through modeling as Markov Decision Process (MDP), in which the primary outcome of interest is an optimal inventory control policy leading to maximum profit.

HICSS

2. Literature review

The nature of our research crosses multiple streams of literature; from perishable inventory management to RFID monitoring of transportation conditions. This study involves and applies knowledge from various disciplines; including biological food science, agribusiness, information technology, simulation programming, decision support, and management science. The following three subsections briefly touch on past and recent contributions that serve as the main branches most closely related to our efforts here.

2.1. Perishable inventory management

A literature review by Karaesmen et al. (2008) referenced over one hundred published papers on research relating to the topic of managing perishable inventory [6]. They classify the literature into fixed and random life time, as well as periodic and continuous review of inventory control policy. While extensive amount of literature on managing perishables with fixed life times can be found, those that focus on random life times can mostly be traced back to Nahmias (1977) [7]. Our study differs from existing ones mainly in that we consider the effects of temperature abuse during transportation, which consequently impacts the remaining life of items received in replenishment. For example, products that arrive today may perish before products received in the previous day.

Recent work by Ketzenberg et al. (2017) focuses on determining a retailer's optimal order quantities and expiration dates for perishable products where perishing is unobservable [8]. While their problem is framed with a cost minimization objective, we assume that perishing is observable and we allow for recapture of lost sales. Thus, we proceed to explore the potential benefits from recapturing unsatisfied demand, along with the impact of overnight replenishment on a retailer's expected net earnings. As a result, our models are framed with profit maximization objectives; taking into account proceeds from units sold.

2.2. Time and temperature monitoring

The value proposed in this work is based primarily on the information obtained from time and temperature history (TTH). Several studies have explored various conditions that perishable products are exposed to at differing parts of the supply chain during shipment [9]. Most fruits and vegetables first enter the supply chain after harvest at ambient temperature. They are often delivered to, or picked up by, regional distribution centers before entering a temperature-controlled chain.

External factors such as distance, time spent, or weather conditions could all cause variations to the amount of remaining life of products. Nunes et al. (2006) report that temperature is the main characteristic of distribution environment to cause the greatest negative impact on shelf life of perishables [10]. An integrated framework for applying RFID monitoring to perishable inventory management can be found, for instance, in Chande et al. (2005) [11].

2.3. Value of information

Sahin and Robinson (2002) as well as Huang et al. (2003) provide broad overview of literature on value of information (VOI) for inventory management [12, 13]. Unlike these contributions, we are most interested in the daily decisions of placing orders for replenishment in a grocery retail setting; given discrete time stochastic demand under periodic inventory review. Therefore we set our views based on the construct of Markov Decision Processes (see, e.g., Puterman (1994)), in which we can analyze the problem in states of inventory and allow the store to take actions periodically by placing orders [14].

Earlier works by Aggoun et al. (1997), and later (1999), establish an integer-valued inventory model for perishable items along with various parameter estimators to find optimal replenishment schedule; in which we are borrowing some notation conventions [15, 16]. Kouki et al. (2010) and (2015) further explore the use of transition probabilities and stationary distributions to estimate impacts of life time variability on cost performance; taking into account lost sale and outdating cost parameters [17, 18].

Studies on VOI gained through implementing RFID in perishable inventory management by Ketzenberg et al. (2015) and dynamic expiration dates by Gaukler et al. (2017) closely resemble our starting point. They have reported up to 43.2% and 41.2%, respectively, in per-period cost reductions on average with no backlogging and one day lead time [19, 20]. It is from here that we continue to expand the literature.

3. Study context and descriptions

Our first study focuses on reducing the impact of lost sales through converting portions of unsatisfied demand by offering customers a choice to backorder stocked-out products. Currently, when a given shopper faces a stock out event, he has to decide whether to purchase an alternative item, go through checkout and buy the particular item elsewhere, or completely drop the cart and take his business to another store. What if we could reduce the impact of a lost sale event by adding flexibility to the system?

The increasing proliferation of smartphone APP usage may one day reach the point where more shoppers tap their phones to pay for groceries instead of searching their wallets for club cards. By having the membership information uniquely tied to each phone, recapturing a lost sale literally becomes just a phone tap away. While not all shoppers will be willing to return next day for pickup, some may not mind having items delivered to their homes. Although stock out events inevitably take place in the presence of outdating costs, the demand itself could possibly be recaptured if implemented in a convenient way to the customer. Incentives such as small discounts or store points can be employed to encourage the success rate of recapturing an unsatisfied demand. It is not clear at this point, though, how much and in what form the incentives would best translate back into increased returns. Perhaps the recapture feature would be useful for a new product during introductory phase where demand is often uncertain.

The second study, closely linked to the first, explores what we term "just-in-time order placement". Grocery retailers with access to TTH information during transport within the distribution chain possess much better abilities to predict the remaining shelf lives of perishable products in their inventory. Knowing how many units of inventory will expire today is useful, and so as knowing if they will tomorrow. Could there remain untapped benefits from insights into the future shelf lives of inventory, in terms of decision making?

Given that the store manager knows when a certain amount of her perishable goods will likely deteriorate beyond top shelf quality, her problem reduces to the two unknowns of demand variability and remaining life of incoming replenishment. With more details in section 4.3, we consider a Markov chain where probabilistic events occur after an action has been taken. It is possible that, by knowing future inventory status, we could make decisions into the future beforehand. And if the action taken can influence expected outcomes within the time frame that the said decision has been made in advance, the manager may be able to improve her store's effectiveness in adapting to the variability of demand.

Since an optimal policy model exists for the one-day lead time period, we adapt it to reflect the ability to influence current period's outcomes. The study could be informative to those already operating with no effective lead times, or others with the ability to increase supply responsiveness and are considering shifting to just-in-time order placements. Such ordering procedures could be quite compatible to a recapture feature. Backlogged units behave as a pull from next period's inventory pool, therefore a more responsive system should provide superior accommodation. Whether or not delaying the order process would be worth putting additional pressure on supplier's end remains to be seen, however.

4. Models

We cast the problem as a Markov Decision Process in a similar manner as the Ketzenberg et al. (2015) and Gaukler et al. (2017) studies [19, 20]. The setting is described from the perspective of a grocery retailer selling perishable food products. The store places an order once a day and receives replenishment from an external supplier. The supplier has adopted RFID monitoring of TTH, and can in turn provide accurate estimates to when a given lot of goods will perish.

Let a stand for the remaining shelf life (in days) of a particular lot of perishable inventory after it has been received by the store. Here we explore a specific case of highly-perishable products having maximum remaining life times on the shelf, M, of 3 days at the time of arrival. Examples include: fresh dairy products, sashimi-grade raw fish, strawberries, ripened tomatoes, and fresh basil leaves. The spreading of age class random variable a, at the time each lot of replenishment is received, can be described by a discrete probability distribution $\varphi(a)$. Each passing day the particular lot of goods remains in inventory, its age class reduces by one. Once a reaches 0 the lot is considered outdated and will be subjected to heavily discounted sale or discarded for a small fee.

Demand d is modeled as discrete, stochastic, and stationary over time, with a mean u_d, probability mass function $\phi(d)$, and coefficient of variation CV_d. As an example, $\phi(d) = \{0.1, 0.4, 0.5\}$ implies that there is a 40% probability that demand for the period will be 1 unit, and so on. The order of events in each period consists of: (i) receive replenishment from order made in the previous period, (ii) allocate the replenishment into stocks of correct age categories and place an order if necessary, (iii) face incoming demand throughout the selling day (period), (iv) reduce age classes of all unsold inventory at the end of the period by 1 and outdate any perished units ($a = 0$) from inventory.

The decision of interest is the quantity of replenishment, q, to order. Therefore we view the system as being in one of a number of possible states (S) of inventory. Let i_a (occasionally written as i_x or i_j) denote the amount of inventory on hand having age class a, and y indicate the maximum space to hold inventory of a particular age class. For example, if $M = 2$ and $y = 1$, we have 4 possible inventory states; namely $\{0,0\}$, $\{1,0\}$, $\{0,1\}$ and $\{1,1\}$. The probability of moving from one state (S) to another (S') by taking action (q) is represented by $P(S : q : S')$, or $P(i_1, \dots, i_M : q : i'_1, \dots, i'_M)$. Let matrix $\bar{P}$ represent all elements of the possible combinations of $(S : q : S')$. If state space was the same as above and q could be 0, 1 or 2, then $\bar{P}$ comprises of 48 elements.

We model the expected reward (or cost) of moving from state S to S' through four cost parameters, including: profit from selling each unit of inventory, s, a penalty p for each unit of lost sale, period holding cost per unit h, and c for the combined costs of outdating a unit of expired inventory. When met with a demand of 1 unit for the selling period, if the system was in state $\{i_1 = 2, i_2 = 1\}$ and $\{s, p, c, h\}$ were $\{3, 1, 2, 0.1\}$, then the immediate reward would consist of making a sale of 1 unit, outdating 1 unit of inventory, and holding 1 unit for a resulting period net profit of 0.9. More details can be found in sections 4.2 and 4.3, but for now we collectively call all combinations of the expected reward of being in state S and taking action q, $R(S:q)$, as elements of matrix $\bar{R}$.

4.1. Assumptions

For traceability and performance reasons, we assume there are no shortages in supply, and a product retains constant utility while its remaining shelf life is at least 1. All units received in the same lot of replenishment will expire at the same time due to undergoing the same environmental conditions during transport. Inventory units are sold by first-to-expire, first-out policy (FEFO) based on information available through RFID monitoring.

4.2. Standard model (STD)

As described in section 3, our first model extends from the current literature by allowing for recapture of unsatisfied demand when a lost sale occurs. We introduce an additional state variable backlogged demand denoted b to keep track of the amount of recaptured demand to be fulfilled in the subsequent period. Let z represent the maximum allowable backlogged demand in units, and r stand for the rate of recapturing a given unit of unsatisfied demand. Our probability of moving within the state space now becomes $(i_1, \dots, i_M, b: q: i'_1, \dots, i'_M, b')$, and the size of $\bar{P}$ increases to $((y + 1)^M * (z + 1))^2 * (q_{max} + 1)$.

The probability of moving from state S to state S' is governed by $\varphi(a)$, $\phi(d)$, and the choice of q primarily through the following inventory balancing equation for $(1 \leq x \leq M)$:

$$i'_x = \left[i_{x+1} - \left(d + b - \sum_{j=1}^{x} i_j \right)^+ \right]^+ \{+q \ if \ a = x\} \quad (1)$$

where $(value)^+$ is equivalent to $max\,(0, value)$.

Let u denote the amount of unsatisfied demand when the system is in state S and facing incoming demand d.

$$u = d + b - \sum_{x=1}^{M} i_x \quad (2)$$

Finally, let the chance of landing in the correct backlog state be defined as:

$$\rho = \frac{u!}{b'! (u - b')!} \cdot r^{b'} \cdot (1 - r)^{u - b'} \quad (3)$$

The variable ρ is used when there exists a positive unsatisfied demand from the given choices of d and a to determine the probability of going to new states within the same resulting inventory space, but containing differing backlog information b'. The recapture rate r can be specified from 0 to 1, and for most cases ρ simply follows binomial distribution based on success and failure to recapture each particular unit of unsatisfied demand.

In the special case where more unsatisfied demand u happens to be recaptured than the maximum backlog capacity z, we also account for probabilities of recapturing more than capacity; up to u.

$$\rho = \sum_{k=0}^{u-b'} \frac{u!}{(b' + k)! (u - b' - k)!} \cdot r^{b'+k} \cdot (1 - r)^{u-b'-k} \quad (4)$$

Now, the probability of going from one state to another can be found through checking the inventory transfer function and ρ value. In essence, this is equal to a sum of all possible combinations of a and d that moves S to S'

$$P(i_1, \dots, i_M, b: q: i'_1, \dots, i'_M, b') = \sum_{a} \sum_{d} \phi(d) \cdot \varphi(a) \cdot \rho \quad (5)$$

The reward of taking action q (ordering q units) while in state $(i_1, \dots, i_M, b)$ is calculated as expected reward over all possibilities of incoming demand. Let $I = \sum_{1}^{M} i_x$, and we have the first component accounting for profit made from units sold.

$$s \cdot \min(I, d + b) \quad (6)$$

We can negate the influence of profit when looking only at costs by setting $s = 0$. The second component, lost sale penalty, comes from unsatisfied demand.

$$-p \cdot (d + b - I)^+ \quad (7)$$

Note that while backlogged demand gets priority to be fulfilled once inventory arrives, in rare circumstances if they are not fulfilled such backlogged demand units are treated the same as a lost sale; meaning they can again be recaptured or lost completely. Next, the outdating cost includes expiring inventory that were not expected to be sold within the state's selling duration.

$$-c \cdot (i_1 - d - b)^+ \quad (8)$$

Lastly, holding costs are assessed on all inventory; less any units that are expected to expire or be sold.

$$-h \cdot [I - max(i_1, d + b)]^+ \quad (9)$$

Thus we can now calculate the expected reward of being in state S and taking action q as:

$$R(i_1, \dots, i_M, b: q) = \sum_{d} [(6) + (7) + (8) + (9)] \phi(d) \quad (10)$$

4.3. Just in time model (JIT)

As noted in the latter part of section 3, our second model assumes a shorter window of replenishment lead time. If the retailer places an order at the end of the selling day and supply can arrive by morning, our system effectively becomes one that has no lead time. This is made possible due to the power of information from RFID monitoring, which allows the store to accurately predict the state of each lot of inventory into the future. By the end of the selling day, an RFID-enabled grocery retailer *already* knows, with a high certainty, the state of inventory tomorrow morning.

The order of events in each period now becomes: (i) receive replenishment from order made at the end of the previous period, (ii) allocate the replenishment into stocks of correct age categories, (iii) face incoming demand throughout the selling period, (iv) reduce age classes of all unsold inventory at the end of the period and outdate perishing units from inventory, and finally (v) placing orders if necessary.

Taking advantage of this knowledge time gap and a responsive supply chain, we can then derive new equations for calculating the $\bar{P}$ and $\bar{R}$ matrices of a Markov Decision Process. First the inventory balancing equation now allows for selling of items from replenishment order that has just been placed.

$$i'_x = \left[i_{x+1} - \left(d + b - \sum_{j=1}^{x} i_j \quad \{-q \ if \ a \le x\} \right)^+ \right]^+ \quad (11)$$
$$\{+q \ if \ a = x + 1\}$$

While delivery seems instantaneous in the model, it is actually done overnight in reality during closed hours. The equation for finding unsatisfied demand u needed during calculation also reflects the change:

$$u = d + b - \sum_{x=1}^{M} i_x - q \quad (12)$$

Finally, the four profit and cost components forming the $\bar{R}$ matrix are updated as well. Again let $I = \sum_{1}^{M} i_x$,

$$\mathbf{R}(i_1, \dots, i_M, b\colon q) = \sum_{d} \begin{bmatrix} s \cdot \min(I + q, d + b) \\ -p \cdot (d + b - I - q)^+ \\ + (14) \quad + (15) \end{bmatrix} \phi(d) \quad (13)$$

The outdating cost term (14) now accounts for the probability that replenishment will arrive with age class 1 and also not expected to be sold.

$$-c \cdot [(i_1 - d - b)^+ + \varphi(1) \cdot (q - (d + b - i_1)^+)^+] \quad (14)$$

Holding cost (15) is assessed on all inventory and incoming replenishment; less any units that are expected to expire or be sold.

$$-h \cdot \begin{bmatrix} I + q - \varphi(1) \cdot \max(i_1 + q, d + b) \\ - (1 - \varphi(1)) \cdot \max(i_1, d + b) \end{bmatrix}^+ \quad (15)$$

With these we are at last able to complete our formation of the reward matrix $\bar{R}$ for the JIT model. While a different action q caused no change to the expected immediate reward for the STD model, it now influences the expected reward matrix $\bar{R}$ due to the seemingly instantaneous arrival of replenishment. It may be helpful in visualizing the problem by assuming that the store can make tomorrow morning's decision, with immediate outcome, by the end of today.

4.4. Definitions and notations

For convenience, a list of all notations along with brief definitions can be found in Table 1 below.

Table 1. Complete list of notations

Symbol	Definition
a	Age class of inventory $\{1 \le a \le M\}$
u_a	Mean age class of incoming replenishment
M	Maximum age class
$\varphi(a)$	Probability mass function of a
q	Order quantity (or action taken)
d	Incoming demand $\{0 \le d \le d_{max}\}$
u_d	Mean of incoming demand
CV_d	Coefficient of variation of demand (SD/mean)
$\phi(d)$	Probability mass function of d
$i_a \ or \ i_x$	Inventory in age class a or x
$i'_a \ or \ i'_x$	Inventory in age class a or x of the next state
y	Maximum inventory in each age class
b	Amount of backlogged demand
b'	Amount of backlogged demand of the next state
z	Maximum backlog capacity in units
r	Rate of recapturing unsatisfied demand
u	Unsatisfied demand
ρ	Index for calculating the chance of going to each b'
s	Profit made from 1 unit sold
p	Penalty for losing a sale (dissatisfaction, reputation)
c	Cost of outdating (purchase + disposal - scrap value)
h	Cost of holding a unit of inventory across one period

5. Method and parameters

Up to this point we are able to create a transition probability matrix $\bar{P}$ and reward matrix $\bar{R}$ for each of the two models. This section describes briefly our procedure of using the $\bar{P}$ and $\bar{R}$ matrices to develop further results. We built the simulation models in R-Studio; a free-to-use software. Optimizations were done using the R package MDPtoolbox developed by Chades et al. found on R depository CRAN [21].

In each experiment, the correct $\bar{P}$ and $\bar{R}$ matrices were constructed for all the given choices of q. We then applied a relative value iteration algorithm that seeks to maximize the long run expected profit (or minimize cost). Once an epsilon-optimal (0.0001) policy that suggests how much q to order when the system is in state S has been found, the iteration process stops. From the optimal policy vector the choice of q was decided for $\bar{P}$ and $\bar{R}$. A stationary distribution π of the state space was calculated through a step search by minimizing the

mean square error; such that $\pi = \pi \bar{P}$. Multiplying the vector R (expected reward of being in state S and taking action q) to the stationary distribution π (the average time that the system spends in each state S), we then arrived at the optimal profit (or cost) for the given parameters.

For continuity, experimental parameters used in Gaukler et al. (2017) were carried over [20]. Demand distribution is kept as negative binomial, has a mean u_d of 5 and CV_d of 0.5 and 0.65; although truncated to max of 15 for performance reasons due to the long right tail. Specific to our highly-perishable product context, a maximum age class of $M = 3$ days was used for items received from replenishment. The two sets of distributions, with equal mean u_a of 2 days, used for remaining life of incoming replenishment are shown in Table 2. To account for profit made from sales in our models, underage cost is split into profit s from sales and penalty p for each stock out event (dissatisfaction, loss of reputation). Note that this ratio $s : p$ only affects the profit/cost output, while outputs such as optimal policy or service level would remain unaffected as long as the values sum to the same total underage cost.

To keep the number of full factorial experiments manageable, cases where loss of reputation or goodwill exceeds the loss in profit from making a sale are not considered. Staple, low margin perishable grocery items such as fresh eggs (where there are practically no substitutes, and dissatisfaction could be particularly high in a stock out event) may not be compatible with a recapture approach. Holding cost was kept constant at 1 cent per unit per period across all experiments.

Table 2. Test parameters for all experiments

Models	STD		JIT
$\varphi(a)$	0.20-0.60-0.20		0.33-0.34-0.33
CV_d	0.5		0.65
c	1	4	8
$s + p$	1	6	12
$s : p$ ratio	1 to 1		3 to 1
r	0	0.3	0.8
z	1		3

Varying strengths of the recapture system (r, z) were tested in both the STD and JIT models. In reality the chance of recapturing an unsatisfied demand is likely tied to the incentives being offered such as a free delivery and discount in price. However since we have not explored yet which level of discount would lead to differing rates of recapture, we instead restrict the problem to the recapture rate itself to estimate the costs saved (or profit gained) by certain recapture rates.

The size of the problem is primarily determined by the amount of information each state needs to carry; in our case being the amount of inventory in each age class and the amount of backlogged demand needs to be fulfilled in the next period. Due to hardware restrictions, we are presently able to test the model up to maximum backlog units of 3 while retaining adequate size of inventory state information. The $\bar{P}$ matrix of the largest experiment presented here contains 1,003,976,272 elements.

We conducted 720 experiments outlined by Table 2. There are a total of 10 *cases* between the STD and JIT models; each with varying recapture strength (r, z). Experiments that carry the same parameters ($\varphi(a)$, CV_d, s, p, c) across all cases are referred to as *sets*, in which there are a total of 72 unique combinations; each representing a specific product facing specific demand.

6. Results

6.1. Optimal expected profit

First, we observe the average long run expected profit produced from the suggested optimal policy; taking into account costs of outdating, losing sales, holding inventory, and profit from sales. At first glance profits improve; in varying scale due to differing parameters, across all cases of experiments. To gain a better perspective of the results, improvements are converted into percentages.

Within the standard (STD) model, up to 16.8 percentage gain can be achieved on average by setting recapture rate to 80% and allowing a maximum of 3 backlogged units. However, the median increase in profit for a more moderate 30% recapture rate only amounts to a mere 2.4%. Results from the just in time (JIT) model shows a similar trend; if not somewhat more capable. It is worth noting that improvements in profit are concentrated to experiment sets with high outdating parameter (purchase cost plus disposal less any scrap value). If only sets that outdating (c) is greater than or equal to the combined underage parameters (s) and (p) are considered, then the average profit improvements rise substantially and now rest within the median and 3rd quartile as illustrated in Figures 1 & 2.
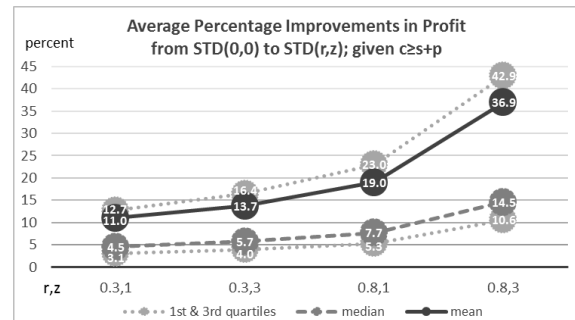

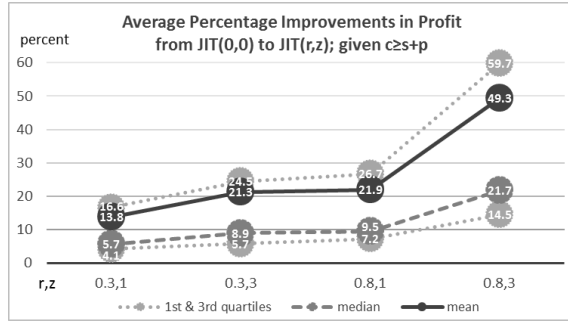Figure 1. Profit improvements within STD cases

Figure 2. Profit improvements within JIT cases

Results from these sets may be of specific interest here since cost related to outdating is the prime distinguishing characteristic of perishable inventory management. Both models, and especially the JIT model, perform increasingly well as the outdating to underage ratio rises. Notably, 6 experiment sets returned unprofitable values at the standard, no recapture case. They were excluded from prior profit improvement analyses, but on average a net loss of 0.32 improves to net profit of 0.38 under the JIT case with ($r = 0.8$, $z = 3$). Though profits remained small, it may become feasible to offer certain products previously not worth stocking. Results after this point cover all 720 experiments.

6.2. Perishable inventory metrics

The average number of units outdated per period reduces considerably as the recapture rate (r) and max backlog slots (z) increase; with the JIT cases consistently starting at lower baselines (Figures 3 & 4).
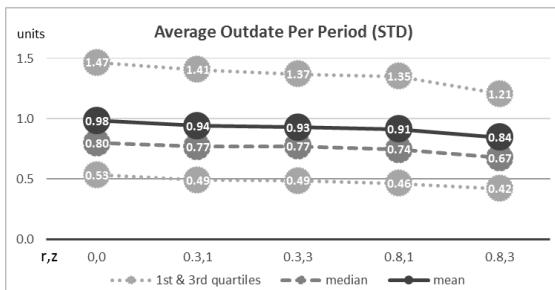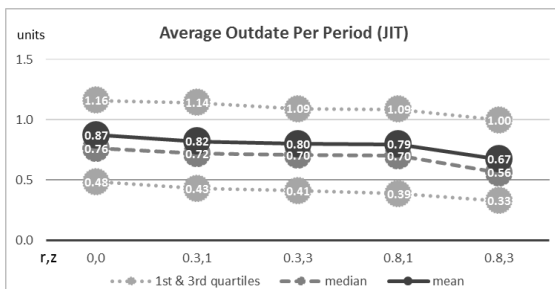

Figure 3. Number of units expected to outdate (STD)

Figure 4. Number of units expected to outdate (JIT)

The flexibility to recapture unsatisfied demand had reduced the overall pressure on the system to maintain high levels of inventory to accommodate occasions of high demand influx. The JIT replenishment model offers additional agility to react to incoming demand; thus reducing stock levels (and outdating) even further.

In terms of freshness, we pay close attention to the remaining life of items at the time the product is sold to customer. Replenishments are received with a mean remaining life of 2 days, as per the study parameters, and continue to deteriorate from there. Figures 5 & 6 show slight improvements in freshness toward ideal conditions; with average values residing right along the median lines.
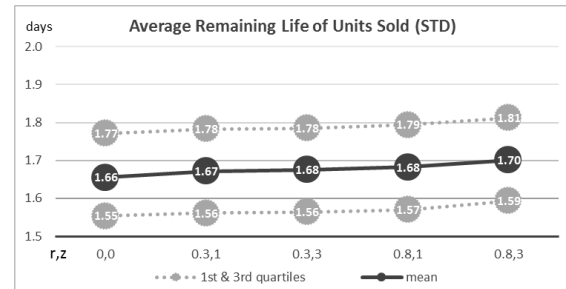

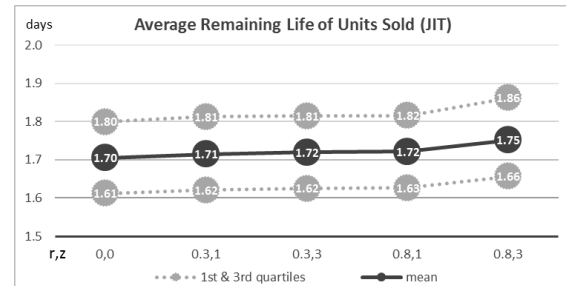Figure 5. Freshness of units expected to be sold (STD)

Figure 6. Freshness of units expected to be sold (JIT)

6.3. Inventory performance metrics

Next, we explore a few operational metrics starting at the average order quantity found in Table 3. As recapture parameters become stronger, average order sizes increase to fulfill additional units of backlogged demand; while at the same time reducing in variability across the quartiles. Unsurprisingly, since the JIT model experiences less outdating events relative to the STD model, it produces slightly lower order quantities as well. The JIT cases offer even tighter quartile ranges, and this narrowing effect may prove to be a welcome unintended-consequence to suppliers upstream.

To gain better understanding of inventory status, we proceed to inspect the average inventory in stock; which differs greatly amongst the two models.

Table 3. Average order quantity and inventory level

Average Order Quantity (units)						
Model		r=0, z=0	0.3, 1	0.3, 3	0.8, 1	0.8, 3
STD	1st qt.	4.76	4.83	4.86	4.86	4.99
STD	mean	<u>5.34</u>	<u>5.37</u>	<u>5.40</u>	<u>5.39</u>	<u>5.49</u>
STD	3rd qt.	6.31	6.26	6.22	6.21	6.14
JIT	1st qt.	4.81	4.91	4.94	4.98	4.95
JIT	mean	<u>5.22</u>	<u>5.27</u>	<u>5.31</u>	<u>5.31</u>	<u>5.36</u>
JIT	3rd qt.	6.02	6.03	5.97	5.98	5.90
Average Inventory in Stock (units)						
STD	1st qt.	6.24	6.22	6.21	6.19	6.16
STD	mean	<u>7.62</u>	<u>7.55</u>	<u>7.55</u>	<u>7.52</u>	<u>7.47</u>
STD	3rd qt.	9.73	9.56	9.46	9.41	8.99
JIT	1st qt.	1.33	1.23	1.22	1.21	0.81
JIT	mean	<u>2.07</u>	<u>1.97</u>	<u>1.92</u>	<u>1.91</u>	<u>1.65</u>
JIT	3rd qt.	2.93	2.89	2.62	2.62	2.41
Average Inventory on Hand (units)						
STD	mean	<u>7.62</u>	<u>7.55</u>	<u>7.55</u>	<u>7.52</u>	<u>7.47</u>
JIT	mean	<u>7.30</u>	<u>7.24</u>	<u>7.23</u>	<u>7.23</u>	<u>7.01</u>

Table 4. Average availability and fill rate

Availability to New Demand (percent)						
Model		r=0, z=0	0.3, 1	0.3, 3	0.8, 1	0.8, 3
STD	1st qt.	90.8	89.6	88.0	88.4	84.5
STD	mean	<u>91.3</u>	<u>90.2</u>	<u>89.1</u>	<u>89.5</u>	<u>85.9</u>
STD	median	<u>95.0</u>	<u>94.0</u>	<u>93.5</u>	<u>93.6</u>	<u>91.2</u>
STD	3rd qt.	98.1	97.8	97.6	97.5	96.6
JIT	1st qt.	92.2	91.6	91.6	91.6	87.3
JIT	mean	<u>92.3</u>	<u>92.0</u>	<u>91.8</u>	<u>91.7</u>	<u>89.4</u>
JIT	median	<u>96.6</u>	<u>95.5</u>	<u>95.2</u>	<u>95.0</u>	<u>93.7</u>
JIT	3rd qt.	98.4	98.3	98.2	98.2	97.3
Average Fill Rate (percent)						
STD	1st qt.	82.9	84.8	85.3	85.5	89.8
STD	mean	<u>86.0</u>	<u>87.5</u>	<u>88.3</u>	<u>88.5</u>	<u>92.0</u>
STD	median	<u>90.3</u>	<u>90.8</u>	<u>91.4</u>	<u>91.7</u>	<u>94.6</u>
STD	3rd qt.	96.2	96.6	96.8	97.1	98.0
JIT	1st qt.	84.5	86.6	88.2	88.5	91.1
JIT	mean	<u>87.0</u>	<u>89.1</u>	<u>90.1</u>	<u>90.4</u>	<u>93.6</u>
JIT	median	<u>92.5</u>	<u>92.2</u>	<u>92.7</u>	<u>93.0</u>	<u>96.1</u>
JIT	3rd qt.	96.5	97.1	97.3	97.7	98.3

While the STD model hovers near 150% of average incoming demand, the JIT model spends most time holding little inventory in stock. For both models, however, the size and variation of inventory levels appear to decrease and narrow, respectively, as the ability to recapture unsatisfied demand increases.

With these information, we can now compare the effective inventory levels across all cases. The latter part of Table 3 suggests that a more responsive retailer is perhaps better positioned to reap the rewards from recapturing unsatisfied demand in terms of periodic inventory level.

Then, we examine the system's performance from the perspective of a potential customer. For simplicity we assume each incoming demand unit represents a given new customer. Despite the trend in benefits presented so far, the tradeoff consequently manifests into a reduction of availability to new demand (Table 4). The stronger the effects of recapturing lost sale, the harder availability compromises. This occurs as the system becomes increasingly reliant on allowing the infrequent surge in demand to go to backlog.

The fact that a backlogged demand is always given priority to fulfill over a new incoming demand further exacerbates the reduction in availability. Thus, care should be exercised when determining cost parameters; especially on the weight of penalty incurred from lost sales. For products with parameters resulting in very low availability, we suggest exploring 'available online-only' with optional in-store pickup.
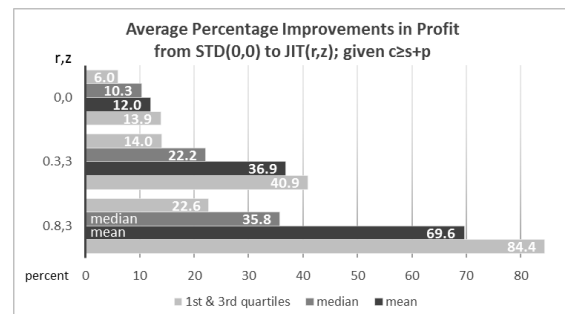
Lastly, the system's overall ability to fulfill demand is summarized as average units sold in each period over mean demand; a type-2 service level. Contrary to the downward trend in availability, Table 4 indicates the reverse is true for fill rate. Sales previously lost without recapture ability may partially cycle back into the system; gaining additional opportunities to be fulfilled.

7. Discussion

Here we package the results to offer an informative view for RFID-enabled retailers currently considering a shift from 1 day lead time to a JIT order placement paradigm; as well as incorporating some degree of recapturing unsatisfied demand. At 30% recapture rate, perishables with high outdating costs may experience up to 22.2% median improvement in net profits relative to the STD case without recapture (Figure 7).

Figure 7. Cross-case profit improvements $\{c \geq s+p\}$

Reduction in number of units expected to outdate per period of 15.7% could be achieved at 30% recapture rate, and at 80% a median of up to 33.5% may be possible (Figure 8). Considering the number of retailers and amount of perishables sold daily, these differences could translate to substantial progress in combating the plaguing issue of food wastes within distribution chain.
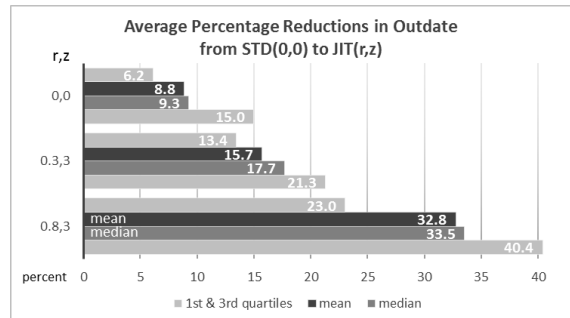

Figure 8. Cross-case reduction in units outdated

Lastly, we address the diverging trends of fill rate and availability. Superimposing the two metrics reveals a more comprehensive guide to service performance (Figure 9). A fresh-groceries retailer may be able to enjoy all the benefits discussed earlier for items with maximum shelf life of 3 days, while minimizing impact on availability, through balancing the recapture rate for desired outcomes. It is worth noting that customers who choose to place backlog orders may receive some compensation as part of the lost sale penalty.
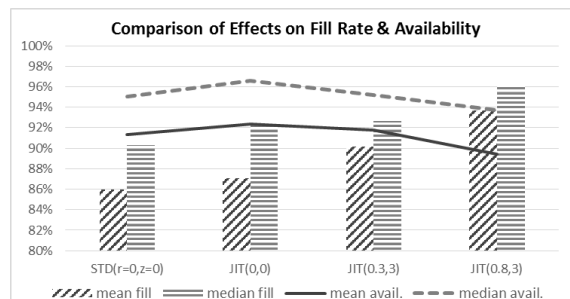

Figure 9. Cross-case guide to fill rate and availability

Given that TTH information is incorporated into a decision support software, the actual implementation of recapture feature simplifies to creating a user-friendly APP that recognizes price tag information, keep track of backlogged demand, and facilitate the fulfilment of merchandise by curbside pickup or home delivery. The improved profit margins could be redistributed back to shoppers in the form of discounts, such as 'get 5% off for curbside pickup, or 'have it delivered free of charge.' Doordash is an emerging food delivery business, where outsourcing can be explored if costs associated with local delivery become prohibitive to bear alone.

A shift from the STD model to JIT may be more complicated, depending on the supplier's ability to respond within a shorter time frame. If the task can be accomplished, though, significant improvements to profit margin, outdating, freshness, and service performance could be attained quite decisively across all areas. When combined, the recapture feature and just-in-time order placement can further extract value from the TTH information obtained through RFID monitoring; potentially expanding the breathing room for an industry operating mostly on slim margins.

On a macro scale, if retailers are able to reduce operating costs, the savings could partially be passed up the value chain all the way back to producers in the form of less pressure on farm prices. When appropriately applied, grocery shoppers can also enjoy fresher products, more selection, as well as better prices.

8. Conclusion

In addition to the various benefits from integrating TTH information in order placing decisions as reported by Ketzenberg et al. (2015), we establish that even more value could be extracted by introducing the recapture feature and just-in-time order placement strategy. Profit improvements of up to 69.6% on average were seen in cases where outdating cost parameters are high. A reduction of nearly one third in the average number of outdated units per period is also reported here. While improvements to freshness remain small, we note that freshness itself was not an optimization parameter. This could present an avenue for research by taking into account, and optimizing for, the freshness parameter.

Intuitively, as the product becomes less perishable, the potential value here decreases. Results are also limited to particularly low maximum shelf life of 3 days, whereas further gains may be achieved under different scenarios. Ultimately, our results provide further evidence of the power of information that RFID and TTH bring to perishables retailing. The ability to accurately predict future states of inventory is shown to positively impact multiple supply chain metrics, including lower inventory levels, smaller order quantities, higher fill rates, as well as better cost performance.

There appear to be opportunities in exploring future research directions. In-depth studies could be done to relate discount factors to recapture rates, improve participation though persuasive design, or expand the MDP models to full omni-channel that adapts to both in-store and online incoming demands; as well as performing cost analyses for differing fulfilment scenarios.

9. Acknowledgements

10. References

[1] Winning with Perishables in a "Value Retailing" World. (2006). White Paper. Park City Group.

[2] Gunders, D. (2012). Wasted: How America is Losing up to 40 Percent of its Food from Farm to Fork to Landfill. Technical report, National Resource Defense Council.

[3] The Future of Grocery: E-commerce, Digital Technology and Changing Shopping Preferences around the World. (2015). The Neilsen Company.

[4] Sands, J. (2016). Sprouts Farmers Market: Assessing Whole Foods Market's 365 Store Rollout. Seeking Alpha.com https://seekingalpha.com/article/3982220.

[5] Hartman, L. (2016). Manufacturers Seeking Natural Ways to Extend Foods Shelf Life. FoodProcessing.com. http://www.foodprocessing.com/articles/2016/natural-ways-to-extend-shelf-life.

[6] Karaesmen, I., Scheller-Wolf, A., Deniz, B. (2008). Managing Perishable and Aging Inventories: Review and Future Research Directions.

[7] Nahmias, S. (1977). On Ordering Inventory When Both Demand and Lifetime are Random. Management Science, 24(1), 82–90.

[8] Ketzenberg, M., Gaukler, G., Salin, V. (2017). Expiration Dates and Order Quantities for Perishables. Technical report, Texas A&M University. Under review by European Journal of Operational Research.

[9] Taoukis, P., Koutsoumanis, K., Nychas, G. (1999). Use of Time-Temperature Integrators and Predictive Modeling for Shelf Life Control of Chilled Fish under Dynamic Storage Condition. Int. J. Food Microbiology.

[10] Nunes, M.C.N., Emond, J.P., Chau, K.V., Rauth, M., Dea, S., Pelletier, W. (2006). Effects of In-store Conditions on Quality of Fresh Fruit and Vegetables. Report to public super markets, University of Florida.

[11] Chande, A., Dhekane, S., Hemachandra, N. (2005) Perishable Inventory Management and Dynamic Pricing Using RFID Technology. Sadhana Vol 30, Parts 2 & 3, June 2015, 445–462.

[12] Sahin, F., Robinson, P. (2002). Flow Coordination and Information Sharing in Supply Chains: Review, Implications, and Directions for Future Research. Decision Sciences, 33(4), 1–32.

[13] Huang, G., Lau, J., Mak, K. (2003). The impacts of Sharing Production Information on Supply Chain Dynamics: A Review of the Literature. International Journal of Production Research, 41(7), 1483–1518.

[14] Puterman, M. (1994). Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, Inc.

[15] Aggoun, L., Benkherouf, L., Tadj, L. (1997). A Hidden Markov Model for an Inventory System with Perishable Items. Journal of Applied Mathematics and Stochastic Analysis, 10:4 (1997), 423-430.

[16] Aggoun, L., Benkherouf, L. (1999). A Stochastic Inventory Model with Perishable and Aging Items. Journal of Applied Mathematics and Stochastic Analysis, 12:1 (1999), 23-29.

[17] Kouki, C., Sahin, E., Jemai, Z., Dallery, Y. (2010). Periodic Review Inventory Policy for Perishables with Random Life Time. 8th International Conference of Modeling and Simulation, MOSIM'10.

[18] Kouki, C., Jouini, O. (2015). On the Effect of Lifetime Variability on the Performance of Inventory Systems. Int. J. of Production Economics, 167 (2015), 23–34.

[19] Ketzenberg, M., Bloemhof, J., Gaukler, G. (2015). Managing Perishables with Time and Temperature History. Production and Operations Management, 24(1).

[20] Gaukler, G., Ketzenberg, M., Salin, V. (2017). Establishing Dynamic Expiration Dates for Perishables: An Application of RFID and Sensor Technology. Int. J. of Production Economics, 193 (2017), 617-632.

[21] Chades, I., Chapron, G., Cros, M.J., Frederick, G., Sabbadin, R. (2017). Markov Decision Processes Toolbox. The Comprehensive R Archive Network.