



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Fixing shelf out-of-stock with signals in point-of-sale data

Howard Hao-Chun Chuang

College of Commerce, National Chengchi University, 64, Sec. 2, ZhiNan Rd., Taipei 11605, ROC, Taiwan

ARTICLE INFO

Article history:

Received 2 December 2016

Accepted 28 October 2017

Available online xxx

Keywords:

Decision support systems

Shelf out-of-stock

Point-of-sale

Audits

Data analytics

ABSTRACT

Shelf out-of-stock (OOS) is a salient problem that causes non-trivial profit loss in retailing. To tackle shelf-OOS that plagues customers, retailers, and suppliers, we develop a decision support model for managers who aim to fix the recurring issue of shelf-OOS through data-driven audits. Specifically, we propose a point-of-sale (POS) data analytics approach and use consecutive zero sales observations in POS data as signals to develop an optimal audit policy. The proposed model considers relevant cost factors, conditional probability of shelf-OOS, and conditional expectation of shelf-OOS duration. We then analyze the impact of relevant cost factors, stochastic transition from non-OOS to OOS, zero sale probability of the underlying demand, managers' perceived OOS likelihood, and even random fixes of shelf-OOS on optimal decisions. We also uncover interesting dynamics between decisions, costs, and probability estimates. After analyzing model behaviors, we perform extensive simulations to validate the economic utility of the proposed data-driven audits, which can be a cost-efficient complement to existing shelf inventory control. We further outline implementation details for the sake of model validation. Particularly, we use Bayesian inference and Markov chain Monte Carlo to develop an estimation framework that ensures all model parameters are empirically grounded. We conclude by articulating practical and theoretical implications of our data-driven audit policy design for retail managers.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

The retail store is the last mile of supply chain management and error-free store execution ensures that the collective efforts of the whole supply chain pay off. Store execution primarily involves moving goods from the backroom to the shelves such that products are available to end consumers (Zondag & Ferrin, 2014). However, in-store logistics are highly labor-extensive and hence store execution is prone to various errors such as shrinkage, misplacement, and faulty transactions (Chuang & Oliva, 2015). Among documented symptoms of poor store execution, shelf out-of-stock (shelf-OOS) is a major problem and refers to the case that an item is in-store (e.g., misplaced or stored in the backroom) but it is unavailable to customers (Papakiriakopoulos, Pramataris, & Doukidis, 2009; Ton & Raman, 2010). The retail giant Walmart recently admitted to a shelf-OOS problem and predicted a \$3 billion opportunity in filling in empty shelves caused by ineffective auditing and re-shelving operations (Dudley, 2014). Walmart even issued an urgent memo that demands store managers to improve grocery performance, which was seriously compromised by non-negligible shelf-OOS ratios (Greenhouse & Tabuchi, 2014).

To solve the shelf-OOS problem that plagues end customers, downstream retailers, and upstream suppliers, store managers of-

ten times ask store employees to perform shelf audits and fill inventory on the shelf. The timing of shelf audits hence has substantial influence on product availability, customer satisfaction, and sales performance (Aastrup & Kotzab, 2010). Therefore, numerous cost-minimization policies have been proposed to assist shelf audit decisions. However, those policies often have one/multiple partially observed state variables (e.g., number of periods or transactions since last audit, shelf inventory level), which may undermine their practical applicability. In this study, we propose a point-of-sale (POS) data analytics approach to shelf audit policy design. Our decision support model is independent of any particular type of shelf inventory replenishment policies and exclusively based on POS data. Being widely available, POS data is reflective of customer demand subject to erroneous store execution. Some retailers have strived to estimate OOS rates from POS data (Gruen & Corsten, 2008). In the age of data analytics, this data-driven modeling choice is deliberate and ensures that the proposed policy is easy-to-implement.

Specifically, we keep track of unlikely events (probabilistic anomalies) in POS or scanned sales data as departures from normal operations and initiate an intervention in a cost-efficient way. In the context of retail operations, z signals, i.e., consecutive zero sales in POS data, are deemed as probabilistic anomalies and strong indicators of shelf-OOS. The use of consecutive zero sales as signals to trigger shelf audits has been proven useful in prior

E-mail address: chuang@nccu.edu.tw<https://doi.org/10.1016/j.ejor.2017.10.059>

0377-2217/© 2017 Elsevier B.V. All rights reserved.

studies (Chuang, Oliva, & Liu, 2016; Fisher & Raman, 2010). Some manufacturers even infer that if a key item has no selling records during a time interval, chances are that the item is not on the retail shelf (Zondag & Ferrin, 2014). In this paper, we develop data-driven models that explicitly account for the fact that an observed zero sale for an item might be attributed to its underlying demand variation or caused by shelf-OOS. Given the realization of z signals, we derive the conditional probability of shelf-OOS and the conditional expectation of shelf-OOS duration. The derivation considers underlying drivers of observed zero sales, accounts for managers' estimate of OOS likelihood, and serves as the core of our data-driven audit policy.

Even though it is possible to include extra state variables such as system inventory records in the model, inventory records in the retail sector are largely erroneous (DeHoratius & Raman, 2008). The proposed policy focuses on maintaining shelf inventory availability and avoids the use of shelf inventory as a state variable, which requires retailers to differentiate shelf inventory records from backroom inventory records. Doing so would increase data requirements and data collection efforts. In fact, data errors are too common to lead to poor decisions as decision makers do not know how much inventory they actually have (Cachon, 2012). Even the item-level RFID does not guarantee full inventory visibility as tag readers (that are imperfect) would result in erroneous shelf inventory records. However, without detecting shelf-OOS caused by execution errors such as shrinkage and misplacement, retailers will not be able to replenish empty shelves and mitigate lost sales even when the backroom capacity is sufficient.

All the afore-mentioned issues create challenges for operations researchers to develop simplistic yet applicable audit policies. Hence, for practical considerations we design the audit policy exclusively based on z signals (i.e., consecutive zero sales) that are distribution-free and fully observable in POS data. After deriving the audit policy, our analysis enables retail managers to better understand the impact of probabilistic factors – stochastic transition from non-OOS to OOS, zero sale probability of underlying demand, and managers' perceived OOS likelihood – on audit decisions. We also uncover interesting interactions between the probabilistic factors and other factors (e.g., sales potential, cost). Further, we conduct simulation experiments to show the applicability and validate the utility of the POS data analytics approach. Extensive simulations uncover scenarios where data-driven shelf audits can effectively improve cost performance and complement existing shelf control systems that are prone to unobservable shelf stock loss. Particularly, simulation results suggest that when demand (and hence shrinkage) rates increase, shelf audits driven by z signals enable retailers to achieve significantly lower system costs.

While our paper is not the first in the literature that proposes decision support models for retail shelf audits, our study contributes to retail operations in two major ways. First, in line with Fisher and Raman (2010) and Chuang et al. (2016), we use the number of consecutive periods of zero sales as the state variable in our model. We improve their work by explicitly accounting for potential causes (i.e., demand variation or shelf-OOS) of realized zero sales, relevant cost factors, and intrinsic sales potential. Moreover, we incorporate the rarely studied random fixes into our model. This non-trivial relaxation of modeling assumptions reveals intricate dynamics of policy behaviors that are attributed to the chance of random fixes and carry important implications for shelf audit decision-making processes. By doing so, we come up with a probabilistic audit policy that is cost-sensitive and more comprehensive. Second, our POS data analytics approach avoids peculiar assumptions and utilizes scanned sales observations that are readily available to retailers. We further develop estimation techniques for key model parameters using maximum likelihood approaches and Bayesian inference with Markov chain Monte Carlo methods.

The nature of Bayesian update also addresses the potentially non-stationary transition matrices in our models. Unlike studies that propose decision support models without showing how to estimate model parameters, our estimation framework ensures that audit decisions are empirically grounded.

The rest of the paper is organized as follows. Section 2 provides a succinct summary of relevant literature. Section 3 presents the design of a POS data-driven shelf audit policy and analyzes behaviors of the proposed policy. Section 4 incorporates random fixes of shelf-OOS into our policy design and sheds light on the impact of random fixes on audit decisions. Section 5 reports a simulation study that quantifies the cost-effectiveness of the proposed model. Section 6 presents primary tasks involved in model validation and estimation. We conclude by discussing key implications of our modeling effort.

2. Related literature

Inventory audits are deemed effective for elevating inventory integrity and product availability, both of which lead to better services and sales (Chuang et al., 2016). Prior studies have developed various *cost-minimization* decision support models under different types of inventory operations. Despite their differences in assumptions and settings, the common objective of those models is to determine the *optimal timing* of inventory audits. Given a re-stocking policy, Iglehart and Morey (1972) propose a cycle-count model that determines frequency and depth of inventory audits. In a similar vein, Kumar and Arora (1992) and Sandoh and Shimamoto (2001), develop models to find optimal frequencies of stock audits based on exponential inter-arrival time of inventory errors. Moving beyond optimal cycle-counting, Kok and Shang (2007) propose a joint audit and replenishment policy. While afore-mentioned studies model errors that cause OOS as random variables, DeHoratius, Mersereau, and Schrage (2008) take a step forward and apply Bayesian inference to construct probability distributions of inventory level. Using Bayesian inventory records, they develop an inventory audit policy based on expected value of perfect information. Different from above studies on the timing of *internal* audits (from retailers' perspectives), Chuang (2015) develops a periodic inventory audit policy for *external* service providers, who (unlike retailers) have limited/no access to inventory/sales information except audit reports. Quantifying the impact of unobserved human errors on optimal audit timing also distinguishes his model from others.

Nearly all foregoing studies on audit policy design do not differentiate store-OOS from the focal issue shelf-OOS in our study. Store-OOS (i.e., zero inventory holdings in both backroom and shelf) requires placing orderings to upstream suppliers, whereas shelf-OOS is more related to in-store logistics (Chuang et al., 2016). Also, prior literature tends to view inventory level as a whole. However, in the retailing sector, store inventories for SKUs are typically composed of backroom and shelf inventories. Most retailers keep track of their inventories at the store level (i.e., the sum of backroom and shelf), but do not know the exact amount of items on the shelf (Condea, Thiesse, & Fleisch, 2012). Consequently, it is common for "freezing" to take place in error-prone store operations (Kang & Gershwin, 2005). That is, even though retailers have abundant backroom inventory, they fail to detect shelf-OOS based on z signals and fill empty shelves before the next auditing and shelving. Our model is unique in that it is designed as a complement to retailers' existing shelf inventory audit and replenishment (rather than store-level inventory governed by automatic store replenishment). The POS data-driven audit initiatives are aimed to mitigate the "freezing", such that on-shelf availability can be maximized and lost sales can be reduced.

Table 1
Variables and parameters.

z	Consecutive periods of zero sales	p	Probability of a zero sale
OOS	An indicator of shelf-OOS	$\sim OOS$	An indicator of non-shelf-OOS
k	Cost per audit	M	Expected per period margin
m	Per unit margin	μ^+	Mean of zero-truncated demand
w	Transition probability into OOS	r	Transition probability of staying OOS
$P(OOS)$	Estimate of shelf-OOS likelihood	τ	Cycle length of shelf inventory control
s	Re-shelving inventory level	S	Targeted shelf inventory level
D	I.I.D. discrete demand in simulation	λ	Mean Poisson demand in simulation
δ	Shrinkage rate as % of demand rate	T	Period length of simulation

Few studies in the literature also have the specific focus on shelf-OOS and propose different models to address this issue. Papakiriakopoulos et al. (2009) adopt machine learning techniques to develop heuristic rules for detecting shelf-OOS. While effective, their methods use more than 10 state variables and data requirements are much higher than our approach. Fisher and Raman (2010) propose a quality control type policy exclusively driven by POS data. Specifically, store managers just need to monitor the event of consecutive zero sales and trigger shelf audits for an items with small probabilities of event occurrence. This simple audit policy based on z signals has been proven effective in reducing shelf-OOS through field experiments (Chuang et al., 2016). Our model that considers extra stochastic elements is a more sophisticated extension of the Fisher and Raman (2010) policy. Moreover, we explicitly incorporate into our model cost factors, which are absent in the two studies (Fisher & Raman, 2010; Papakiriakopoulos et al., 2009) but definitely critical to shelf audit decision-making.

3. Policy design and analysis

3.1. Formulation

We employ z signals to develop a data-driven shelf audit policy, since runs of zero sales (i.e., the so-called “freezing” scenario) are shown to be strong indicators of shelf-OOS in both theoretical (Chuang & Oliva, 2015; Kang & Gershwin, 2005) and field studies (Chuang et al., 2016; Fisher & Raman, 2010). Table 1 summarizes key model variables and parameters.

To begin with, we define z as the number of consecutive periods of zero sales. However, the observed event z is not necessarily a faithful representation of customer demand due to the unobserved status of a SKU – OOS (i.e., empty shelves) or $\sim OOS$. The fact that the realized z is a censored observation makes interpreting the z signals difficult and creates an extra layer of complexity for audit policy design.

The OOS can be viewed as a latent variable that reflects the underlying status of a SKU. Under a discrete i.i.d. demand process (e.g., Poisson, negative binomial), which is common in retailing (Chuang et al., 2016; DeHoratius et al., 2008), we can consider the underlying demand for an item in each period (e.g., hour, day) as a Bernoulli trial with a probability p of a zero sale. First of all, we derive the likelihood of an OOS conditioning on z consecutive zero sales, namely $P(OOS|z)$. The derivation involves state transition across z periods. Hence, we need to further define a transition matrix

<i>status</i>	OOS	$\sim OOS$
OOS	1	0
$\sim OOS$	w	$1 - w$

where the parameter w is used to capture the stochastic transition from $\sim OOS$ into OOS. The matrix is essentially a two-state Markov chain that governs the status of a SKU between periods. The parameter w captures store execution equality and is expected to be higher when transaction errors and shrinkage rates are

common. The level of w can also be used to adjust perceived risks of shelf-OOS. Note that w could be modeled as a function of other state variables such as shelf space and time-variant inventory records. However, doing so complicates the model formulation but also increases data requirements. State variables like inventory records also typically contain errors (Chuang et al., 2016). Hence, we deliberately make the transition probability an estimable constant for the ease of model formulation and application. We further note that the transition matrix could be non-stationary (i.e., w may drift over time). In Section 6 we address the issue of non-stationarity and develop a Bayesian statistical model that allows decision-makers to continuously update parameters.

By the end of each period, the POS record for an SKU is observed by a store manager who can assess $P(OOS|z)$ that is derived as follows.

Proposition 1. Under a discrete i.i.d. demand process and given that z consecutive zero sales are observed

$$P(OOS|z) = \frac{P(OOS)w^{\frac{1-[(1-w)p]^z}{1-(1-w)p}}}{P(OOS)w^{\frac{1-[(1-w)p]^z}{1-(1-w)p}} + (1 - P(OOS))(1 - w)^z p^z}$$

Proof. Through mathematical induction, it is easy to show that

$$P(z|OOS) = w \sum_{i=0}^{z-1} [(1-w)p]^i = w \frac{1 - [(1-w)p]^z}{1 - (1-w)p}$$

$$P(z | \sim OOS) = [(1-w)p]^z$$

Hence,

$$P(z) = P(OOS) \cdot P(z|OOS) + P(\sim OOS) \cdot P(z | \sim OOS) \\ = P(OOS)w \frac{1 - [(1-w)p]^z}{1 - (1-w)p} + (1 - P(OOS))(1-w)p^z$$

$$\text{Finally, } P(OOS|z) = \frac{P(OOS)w^{\frac{1-[(1-w)p]^z}{1-(1-w)p}}}{P(OOS)w^{\frac{1-[(1-w)p]^z}{1-(1-w)p}} + (1-P(OOS))(1-w)^z p^z} \quad \blacksquare$$

$P(OOS)$ in Proposition 1 is a prior probability estimate that characterizes a store manager’s expectation for any SKU to exhibit shelf-OOS. Note that $P(OOS)$ and w can be used to capture store managers’ perceived risk of lost sales. Managers could reflect their subjective beliefs about the odds of shelf-OOS by adjusting the two probabilistic parameters. In Section 6 we will introduce a Bayesian statistical model to formally infer those parameters from data. We further derive

Proposition 2. Under a discrete i.i.d. demand process, the expected number of shelf-OOS periods for a SKU, when z consecutive zero sales are observed and the item is in the OOS state is

$$\frac{p(w-1)}{1+p(w-1)} - \frac{z}{(p-pw)^z - 1}$$

Proof. Let T denote a random number of shelf-OOS periods when z consecutive zero sales are observed and the item is in OOS state. The possible value of T lies in $[1, z]$. Based on the two-state Markov chain defined earlier, we know that

$$P(T = 1) = \frac{w((1-w) \cdot p)^{z-1}}{\sum_{t=1}^z w \cdot ((1-w) \cdot p)^{z-t}}$$

$$P(T = 2) \frac{w((1-w) \cdot p)^{z-2}}{\sum_{t=1}^z w \cdot ((1-w) \cdot p)^{z-t}}$$

...

$$P(T = z) \frac{w}{\sum_{t=1}^z w \cdot ((1-w) \cdot p)^{z-t}}$$

where the denominator ensures that probability densities are properly scaled and sum to one. By definition $E[T]$ is

$$\sum_{t=1}^z t \cdot P(T = t) = \sum_{t=1}^z t \cdot \frac{w \cdot ((1-w) \cdot p)^{z-t}}{\sum_{t=1}^z w \cdot ((1-w) \cdot p)^{z-t}}$$

For $j < 1$, $\sum_{t=1}^z j^{z-t} = \frac{j^z - 1}{j - 1}$. So the expression above reduces to $\frac{p(w-1)}{1+p(w-1)} - \frac{z}{(p-pw)^z - 1}$. ■

Propositions 1 and 2 serve as the core of our data-driven shelf audit policy design. Note that the two propositions rely on p estimates from demand observations and hence hold for any i.i.d. discrete demand distributions. Let k be the cost of conducting a shelf audit and M denote the expected per-period retail margin for a particular SKU. Following the principle of balancing cost with and without auditing (O'Reagan, 1969), our policy is defined at a break-even point where the cost of auditing (k) is less than or equal to the expected cost of leaving potential shelf-OOS unfixed. Essentially the optimal threshold hinges on "auditing this period" versus "auditing next period", and activates once the potential loss of waiting for an additional period becomes large enough. Specifically, z^* is

$$z^* = \min\{s \in (1, 2, 3, \dots) : k \leq M \cdot P(OOS|s) \cdot \left(\frac{p(w-1)}{1+p(w-1)} - \frac{s}{(p-pw)^s - 1} \right)\} \quad (1)$$

where z^* refers to the number of consecutive zero sales a decision-maker should tolerate before triggering an audit, and $P(OOS|s)$ is shown in Proposition 1. In the RHS of the inequality we multiply the expected per-period margin M by both $P(OOS|z)$ and expected number of shelf-OOS periods (derived in Proposition 2) in order to estimate expected profit loss over the course of event z . However, M is a function of the zero sale probability p and other factors. Specifically, M is the product of per unit margin m and expected per period demand μ . The latter (μ) can be written as $\mu^+ \cdot (1-p)$ where μ^+ is the expected zero-truncated demand per period. The μ^+ can be directly estimated from non-zero POS observations (regardless of assumptions regarding demand distributions) and reflects the underlying sales potential of this SKU. As such, the expected per-period margin $M = m \cdot \mu^+ \cdot (1-p)$ explicitly depends on p and the optimality condition in (1) can be re-written as

$$z^* = \min\left\{s \in (1, 2, 3, \dots) : \frac{k}{m} \leq \mu^+(1-p) \cdot P(OOS|s) \cdot \left(\frac{p(w-1)}{1+p(w-1)} - \frac{s}{(p-pw)^s - 1} \right)\right\} \quad (2)$$

The proposed audit policy in Eq. (2) explicitly strikes the balance between cost of conducting an audit and cost of leaving potential shelf-OOS unfixed. The cost-balancing ideal of our model also matches the need of store managers who launch shelf audits on a continuous basis due to the recurring nature of shelf-OOS.

3.2. Evaluation

As shown in Eq. (2), z^* is a function of cost factors (k, m), demand components (p, μ^+), and probability estimates ($w, P(OOS)$). Proposition 3 first uncovers how the cost factors and sales potential affect policy behaviors.

Proposition 3. The optimal period of consecutive zero sales that triggers an audit (z^*) is

- (a) non-decreasing in the audit cost k
- (b) non-increasing in the per-period margin m
- (c) non-increasing in the expected zero-truncated demand μ^+

Proof. The optimality condition in Eq. (2) can be re-written as

$$\frac{k}{m\mu^+(1-p)} \leq \frac{P(OOS)w \frac{1-[(1-w)p]^z}{1-(1-w)p}}{P(OOS)w \frac{1-[(1-w)p]^z}{1-(1-w)p} + (1-P(OOS))(1-w)^z p^z} \cdot \left(\frac{p(w-1)}{1+p(w-1)} - \frac{z}{(p-pw)^z - 1} \right).$$

For the two factors in the RHS of the inequality, it is easy to show that

$$\frac{d}{dz} \left(\frac{P(OOS)w \frac{1-[(1-w)p]^z}{1-(1-w)p}}{P(OOS)w \frac{1-[(1-w)p]^z}{1-(1-w)p} + (1-P(OOS))(1-w)^z p^z} \right) \geq 0 \text{ and}$$

$$\frac{d}{dz} \left(\frac{p(w-1)}{1+p(w-1)} - \frac{z}{(p-pw)^z - 1} \right) \geq 0$$

Ceteris paribus, when k in the LHS of the inequality increases, z in the RHS needs to be non-decreasing such that the RHS $\geq k$. Similarly, when m/μ^+ in the LHS increases, z in the RHS could be lower or constant (i.e., non-increasing) such that the inequality is still satisfied. Thus, Propositions 3(a), (b), and (c) are proved. ■

Proposition 3(a) shows that a high audit cost (k) leads to a more stringent threshold for triggering shelf audits, whereas Proposition 3(b) indicates that a high profit loss (m) results in a less stringent triggering threshold. Proposition 3(c) implies that audit initiatives tend to be more aggressive when the intrinsic sales potential for the SKU is higher (a larger μ^+). The audit policy behaves in an economically sensible way. The impact of probability estimates ($p, P(OOS), w$) – on z^* , however, has limited analytical tractability. Instead, we perform a series of numerical studies and come up with the observation on policy responses to those probabilistic factors.

Observation. The optimal period of consecutive zero sales that triggers an audit (z^*) tends to be

- (a) non-decreasing in the intrinsic zero sale probability p
- (b) non-increasing in the perceived OOS likelihood ex ante $P(OOS)$
- (c) non-increasing in the ~OOS to OOS transition probability w

Observation (a) implies that an intrinsically high zero sale probability p makes z signals look less like probabilistic anomalies and tends to result in a more stringent triggering threshold. Observation (b) suggests that high shelf-OOS likelihood ex ante ($P(OOS)$) is likely to lead to a less stringent triggering threshold. Lastly, Observation(c) implies that when an item is vulnerable to execution errors and more likely to turn into OOS (i.e., high w), the triggering threshold would be less stringent.

In addition to the first-order effects of model parameters on z^* mentioned in Proposition 3 and Observation above, Figs. 1 and 2 below show some interesting dynamics among model parameters. Given fixed estimates of $P(OOS)$ and w , Fig. 1 reveals interaction effects between k/m and p under low and high levels of sales potential (μ^+). As shown in the left panel, when p is low (e.g., $p = 0.1$), zero sales are probabilistic anomalies and hence z^* remains low while being insensitive to increase in audit cost (k/m). However, when the chance of zero sales is intrinsically high (e.g., $p = 0.8$), z^* remains high and increases with k/m quickly. Consistent with Proposition 3(c), the right panel shows that z^* is overall

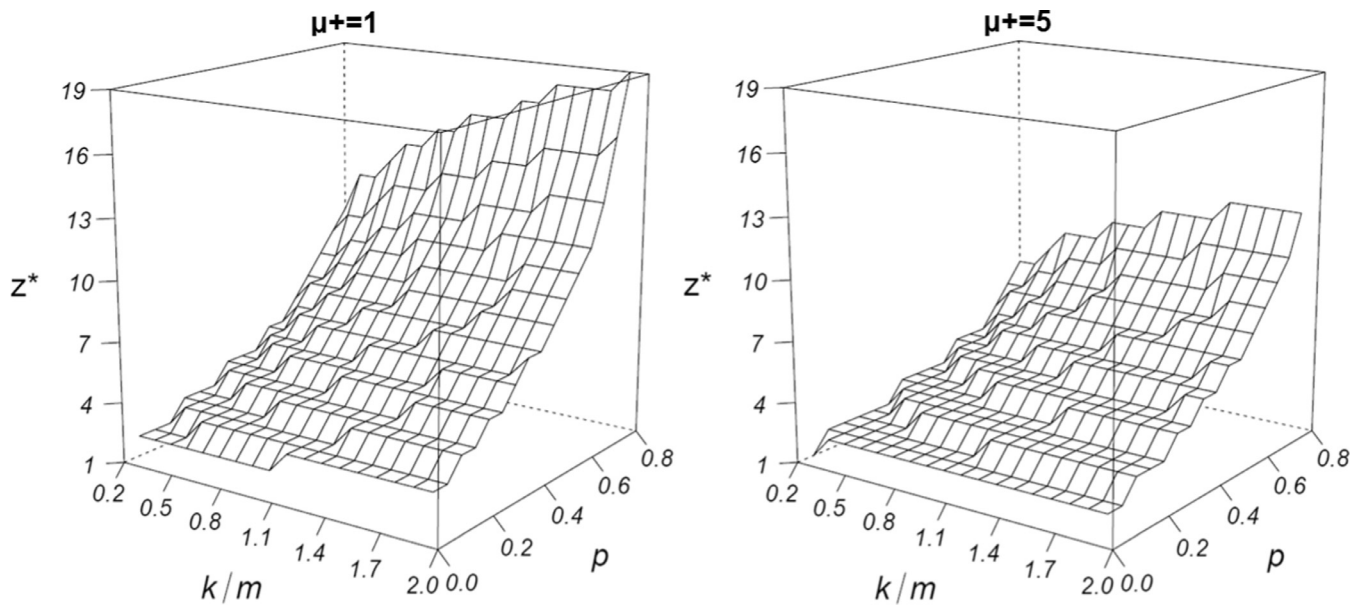


Fig. 1. Effects of k/m and p on z^* ($P(OOS)=w = 0.05$).

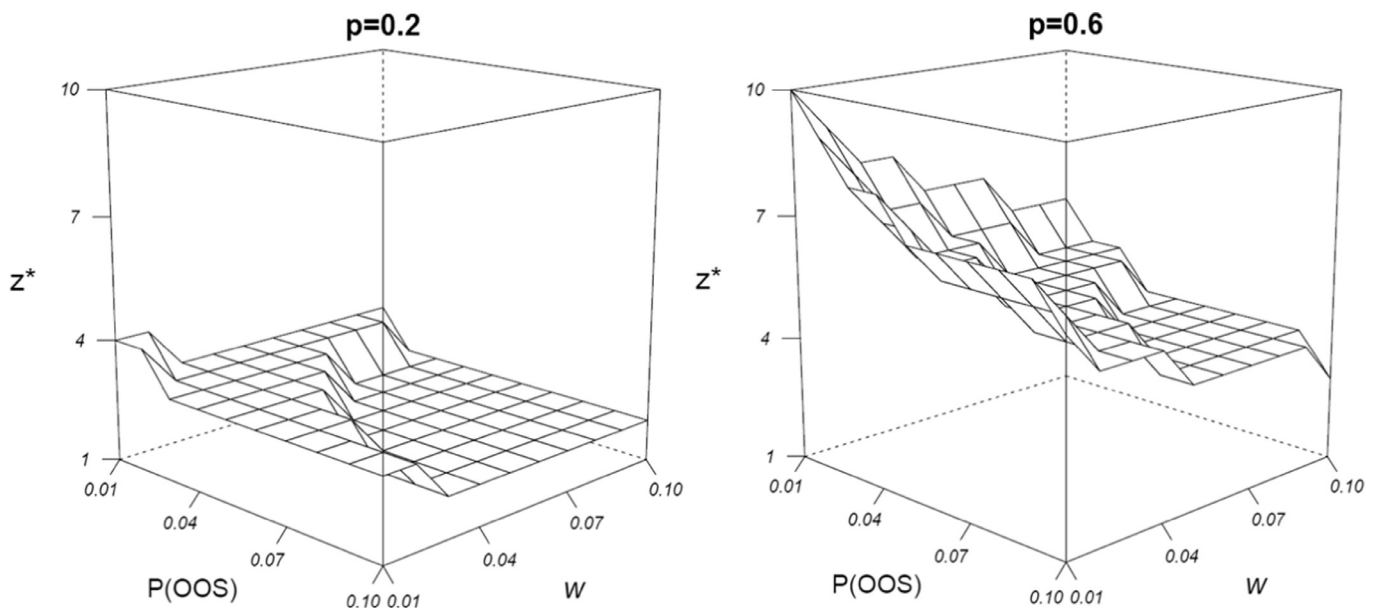


Fig. 2. Effects of $P(OOS)$ and w on z^* ($k/m = 0.5$ and $\mu^+ = 5$).

much lower given a higher μ^+ . That said, we still observe similar second-order interaction effects between k/m and p described above.

Fig. 2 shows effects of $P(OOS)$ and w on z^* under fixed k/m and μ^+ . The left panel shows weak interactions between $P(OOS)$ and w under relatively small zero sale probability ($p = 0.2$). The right panel shows stronger effects of $P(OOS)$ and w on z^* . Consistent with Observations (b) and (c), here z^* tends to decrease given increases in $P(OOS)$ and w . Also, compared to z^* in the left panel, z^* in the right panel is significantly larger due to higher zero sale probability ($p = 0.6$). A major takeaway is that the effect of $P(OOS)/w$ on z^* is largely moderated by p . Taken together, Figs. 1 and 2 reveal substantial interaction effects among cost factors and probability estimates, all of which need to be considered in decision-making processes of shelf audits.

4. Effect of random fixes

4.1. Formulation

The foregoing analysis assumes that once a SKU falls into the OOS status, it will not be fixed before store management triggers a shelf audit and replenishment. However, store employees could randomly walk through the aisles and fix a few shelf stock-outs (Nachtmann, Waller, & Rieske, 2010). Also, some product manufacturers may ask external audit associates to step into stores and fix shelf-OOS in order to ensure product availability (Chuang, 2015; Chuang et al., 2016). To capture the effect of such “random fixes”, we modify the Markovian transition matrix that governs the status

of a SKU between periods as

$$\begin{matrix} \text{status} & \text{OOS} & \sim \text{OOS} \\ \text{OOS} & r & 1-r \\ \sim \text{OOS} & w & 1-w \end{matrix}$$

where OOS is the same variable that reflects the underlying status of a SKU. The parameter r is the probability that a SKU will remain faulty and $1-r$ is the probability that a SKU with empty shelves will be fixed by the end of each period. Based on the updated transition matrix with a possibility of random fixes, we further derive

Lemma 1. Under an i.i.d. discrete demand process with random fixes of shelf-OOS, and given that z consecutive zero sales are observed

$$P(z|\text{OOS}) = P(z-1|\text{OOS})r + P(z-1|\sim \text{OOS})w, \text{ and}$$

$$P(z|\sim \text{OOS}) = P(z-1|\sim \text{OOS})(1-w)p + P(z-1|\text{OOS})(1-r)p$$

$$\forall z \geq 2. P(z=1|\text{OOS}) = w; P(z=1|\sim \text{OOS}) = (1-w)p.$$

For $z \geq 2$, the recursive equations in Lemma 1 can be re-written in a matrix form

$$\begin{bmatrix} P(z|\text{OOS}) \\ P(z|\sim \text{OOS}) \end{bmatrix} = \begin{bmatrix} r & w \\ (1-r)p & (1-w)p \end{bmatrix} \begin{bmatrix} P(z-1|\text{OOS}) \\ P(z-1|\sim \text{OOS}) \end{bmatrix}.$$

Accordingly, we can show that

$$\begin{bmatrix} P(z|\text{OOS}) \\ P(z|\sim \text{OOS}) \end{bmatrix} = \begin{bmatrix} r & w \\ (1-r)p & (1-w)p \end{bmatrix}^{z-1} \begin{bmatrix} w \\ (1-w)p \end{bmatrix}.$$

Given $P(z|\text{OOS})$ and $P(z|\sim \text{OOS})$ from Lemma 1, we can derive $P(\text{OOS}|z)$ with random fixes.

Proposition 4. Under an i.i.d. discrete demand process with random fixes of shelf-OOS, and given that z consecutive zero sales are observed

$$P(\text{OOS}|z) = \frac{P(\text{OOS})P(z|\text{OOS})}{P(\text{OOS})P(z|\text{OOS}) + (1-P(\text{OOS}))P(z|\sim \text{OOS})}$$

where $P(z|\text{OOS})$ and $P(z|\sim \text{OOS})$ are calculated from

$$\begin{bmatrix} P(z|\text{OOS}) \\ P(z|\sim \text{OOS}) \end{bmatrix} = \begin{bmatrix} r & w \\ (1-r)p & (1-w)p \end{bmatrix}^{z-1} \begin{bmatrix} w \\ (1-w)p \end{bmatrix}$$

Proof. Given $P(z|\text{OOS})$ and $P(z|\sim \text{OOS})$ in Lemma 1, this is a direct result of the Bayes' theorem and law of total probability. ■

When considering random fixes of shelf-OOS, the expected number of stock-out periods for a SKU in the OOS state at period z cannot be derived in closed form. Nevertheless, if we consider the stochastic transition for a SKU across different z periods as a binary tree, there will be a total of 2^{z-1} paths that lead to shelf-OOS at period z . Let $j = 1, 2, \dots, 2^{z-1}$ denote each path of the binary tree with probability π_j and $P(z|\text{OOS})$ in Lemma 1 is essentially $\sum_{j=1}^{2^{z-1}} \pi_j$. Also, depending on the number of random fixes in path j , each path has a corresponding length of OOS period $t_j \in [1, z]$. Take $z = 3$ for instance, there will be a total of 4 paths to shelf-OOS at period 3, and each path has an OOS duration length. Specifically,

$$\pi_1 = wr^2, t_1 = 3$$

$$\pi_2 = w(1-r)pw, t_2 = 2$$

$$\pi_3 = (1-w)pwr, t_3 = 2$$

$$\pi_4 = (1-w)^2p^2w, t_4 = 1$$

Similar to Proposition 2, when z consecutive zero sales are observed and the item is in the OOS state at period z we can compute the expected number of shelf-OOS periods for the SKU under probable random fixes as

$$\sum_{j=1}^{2^{z-1}} t_j \cdot \frac{\pi_j}{\sum_{j=1}^{2^{z-1}} \pi_j}$$

where π_j is divided by $\sum_{j=1}^{2^{z-1}} \pi_j$ such that the conditional probability densities are properly scaled and sum to one. Accordingly, we modify the optimality condition of z^* with random fixes as

$$z^* = \min \left\{ s \in (1, 2, 3, \dots) : \frac{k}{m} \leq \mu^+(1-p) \cdot P(\text{OOS}|s) \cdot \left(\sum_{j=1}^{2^{s-1}} t_j \cdot \frac{\pi_j}{\sum_{j=1}^{2^{s-1}} \pi_j} \right) \right\} \quad (3)$$

where $P(\text{OOS}|s)$ is defined in Proposition 4. The optimality condition in Eq. (3) has very limited analytical tractability. Therefore, in the next section we perform a detailed numerical analysis of z^* to better understand the effect of random fixes on probability estimates and audit policy behaviors.

4.2. Evaluation

Since $P(\text{OOS}|z)$ is key to our shelf audit policy, we first assess the impact of random fixes on the shelf-OOS likelihood conditioning on z consecutive zero sales. The left panel of Fig. 3 shows the impact of r on $P(\text{OOS}|z)$ under $p = 0.2$ and $P(\text{OOS}) = w = 0.05$. When random fixes are absent (i.e., $r = 1$), $P(\text{OOS}|z)$ rises to 1 quickly as z increases. However, when there is a 5% ($r = 0.95$) or 10% ($r = 0.9$) chance for an OOS item to be randomly fixed, $P(\text{OOS}|z)$ saturates to a point that increases in r and $P(\text{OOS}|z)$ does not necessarily converge to 1. A somewhat similar pattern of $P(\text{OOS}|z)$ under $p = 0.6$ and $P(\text{OOS}) = w = 0.05$ is shown in the right panel of Fig. 3, where a higher possibility of random fixes (i.e., a smaller r) slows the growth of $P(\text{OOS}|z)$ in observed z .

Given the non-trivial impact of r on $P(\text{OOS}|z)$ shown in Fig. 3, the probability of random fixes $(1-r)$ is expected to affect z^* as well. Fig. 4 shows effects of k/m and p on z^* under low/high r values and fixed levels of $P(\text{OOS})$, w , and μ^+ . While the interaction between k/m , p , and z^* in both panels of Fig. 4 is consistent with the pattern in Fig. 1, it is obvious that random fixes have substantial impacts on optimal timing of triggering shelf audits. Compared to the left panel with 10% chance of random chances, the right panel shows that z^* becomes mostly smaller with only 1% chance of random fixes ($r = 0.99$). This makes sense as the chance of a SKU being fixed before intervention increases (i.e. r decreases), some of the consecutive zero sales observations are more likely to be attributed to demand variation as opposed to shelf-OOS. Hence, decision-makers would be willing to wait longer runs of zero sales (i.e., a larger z^*).

Fig. 5 shows effects of $P(\text{OOS})$ and w on z^* under low/high r values and fixed levels of k/m , p , and μ^+ . Similar to findings above, z^* is overall smaller when the chance of random fixes drops (i.e., $r = 0.98$ in the right panel). Nonetheless, it is interesting to note that the impact of $P(\text{OOS})/w$ seems stronger when the chance of random fixes increases ($r = 0.9$ in the left panel). For instance, z^* reaches its highest value 9 under $P(\text{OOS}) = 0.01$ and $w = 0.01$. This makes sense as in this particular case an item would be less prone to shelf-OOS and thus managers can wait for longer runs of consecutive zero sales. To sum up, Figs. 4 and 5 collectively reveal subtle dynamics introduced by random fixes in the decision-making processes of retail shelf audits.

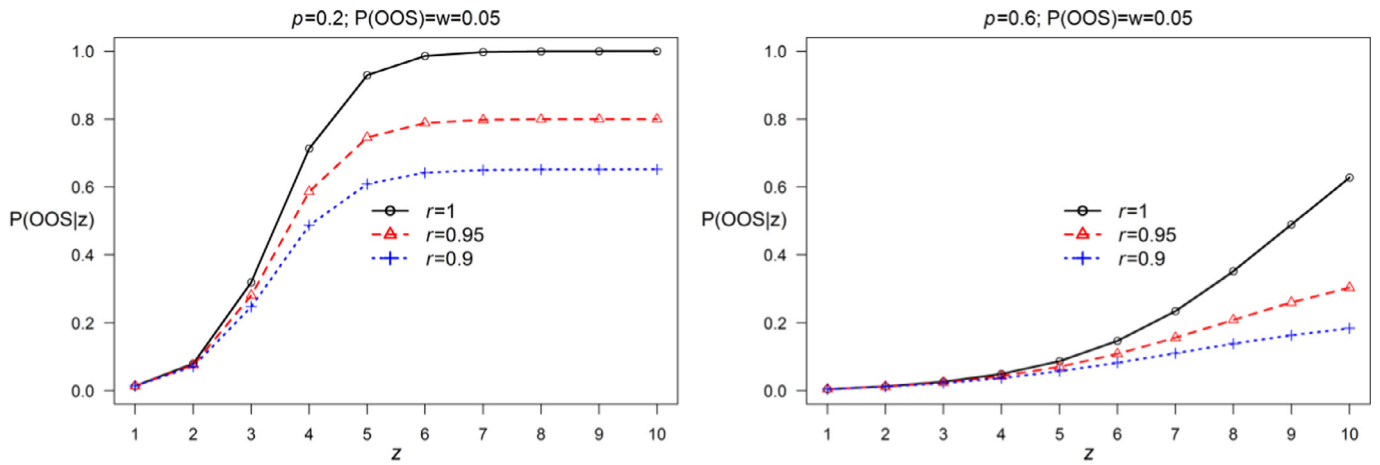


Fig. 3. Effects of random fixes on $P(OOS|z)$.

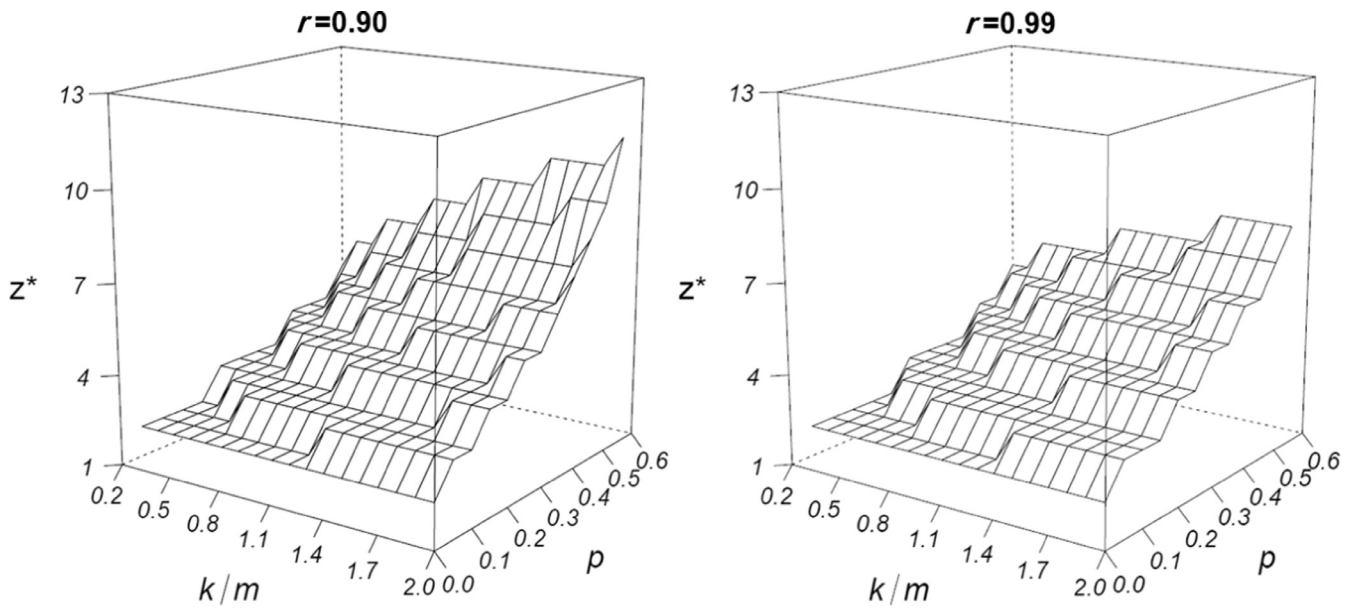


Fig. 4. Effects of k/m , p , and r on z^* ($P(OOS)=w=0.05$ and $\mu^+=3$).




Fig. 5. Effects of $P(OOS)$, w , and r on z^* ($k/m=1$, $p=0.4$ and $\mu^+=3$).

To sum up, the foregoing analysis sheds light on the negative impacts of ignoring potential random fixes of shelf-OOS. Specifically, naively assuming $r = 1$ (i.e., no random fixes) leads to overestimation of $P(OOS|z)$, and consequently, overestimation of the potential profit loss in the RHS of the optimality condition in Eq. (3). As a result, decision-makers tend to take too proactive audit initiatives (i.e., smaller z^*) that may increase wasted audit burdens and create more workload for store associates, who are typically busy with handling in-store logistics and customer transactions.

5. Simulation experiments

To assess the cost benefits of the proposed shelf audit policy, we conduct simulation experiments in which each time period is an hour. With advances in information technologies, deploying the proposed POS data monitoring mechanism on an hourly basis is practically doable. For instance, Fisher and Raman (2010) report a case from Albert Heijn, which uses consecutive hours of zero sales (i.e., z signals) in POS data to detect shelf failures. This time granularity is consistent with retailers that control store inventory for SKUs by days/weeks and control on-shelf inventory by operating hours. In line with prior studies (e.g., Condea et al., 2012) and retailers we work with, we simulate a single-item (τ, s, S) shelf inventory control policy that assumes the backroom to have sufficient supply. Specifically, for a retail store that on average operates τ hours per day, a store associate performs a manual counting before the store opens. That is, a periodic inspection takes place every τ hours, and the shelf has a capacity of S units for this particular SKU. If the on-shelf inventory level reaches or goes below a threshold s during manual counting, the store associate will pick products from the backroom and replenish the shelf. Under the (τ, s, S) shelf operations, unobservable shrinkage (i.e., stock loss) oftentimes occurs and results shelf-OOS within an audit cycle of τ hours. Hence, we expect the proposed data-driven audits based on z signals to complement the shelf operations and enhance on-shelf availability. Below describes the sequence of events in simulating the (τ, s, S) shelf inventory control for $t = 1, \dots, T$ periods, where t is the hour/period index.

1. For t in 1 to T , when t is on the re-shelving cycle τ of (τ, s, S) shelf operations, update the counter of shelf audits $audit(t) = audit(t-1) + 1$. Count the beginning amount of shelf inventory (SI), i.e., beginning $SI(t)$. Replenish the shelf to S and update beginning $SI(t)$ if beginning $SI(t) <= s$. When t is not on the cycle τ , update the counter of shelf audits $audit(t) = audit(t-1)$ and set beginning $SI(t) = ending\ SI(t-1)$.
2. Demand ($D(t)$) for hour t arrives. Given beginning $SI(t)$, sales(t) and shrinkage(t) take place. The sales quantity (sales(t)) in hour t is

$$\begin{cases} D(t) & \text{if } D(t) \leq \text{beginning } SI(t), \\ \text{Round}\left(\text{beginning } SI(t) \frac{D(t)}{D(t)+\text{shrinkage}(t)}\right) & \text{otherwise.} \end{cases}$$

3. Update ending $SI(t) = \text{beginning } SI(t) - \text{sales}(t) - \min(\text{beginning } SI(t) - \text{sales}(t), \text{shrinkage}(t))$. After that compute cumulative lost sales $\text{cum_loss}(t) = (D(t) - \text{sales}(t)) + \text{cum_loss}(t-1)$.
4. When sales(t)=0, update the counter of consecutive zeros sales $z(t)$. If $z(t) >= z^*$, update the counter of shelf audits $audit(t) = audit(t) + 1$ and trigger a shelf audit (reset the count $z(t)$ to 0). Replenish ending $SI(t)$ to S and update ending $SI(t)$ if ending $SI(t) <= s$.
5. After T hours of operations, we compute the cost per period in each simulation run

$$\frac{k \cdot \text{audit}(T) + m \cdot \text{cum_loss}(T)}{T}$$

Step 1 above follows the (τ, s, S) shelf inventory audit and replenishment policy in Condea et al. (2012), whereas steps 2 and 3

follow Kang and Gershwin (2005) who model shrinkage as unobserved sales in inventory systems. Step 4 characterizes extra shelf audit initiatives triggered by z signals, and finally, step 5 returns cost performance. Note that holding costs are negligible and absent in step 5, since backroom inventory is assumed to be sufficient and stable for brevity.

We assume that hourly demand D follows a Poisson distribution with mean λ (Condea et al., 2012; Fisher & Raman, 2010), and shrinkage follows a Poisson distribution with its mean parameter as $\delta\%$ of λ (Kang & Gershwin, 2005). The mean demand parameter λ also determines p and μ^+ for our z^* audit policy. The parameter δ captures different types of invisible shelf stock loss including employee theft, customer theft, spoilage, etc. We set profit margin per unit sold $m = 1$ and perform the simulation for 1000 runs under $\delta = 10\%$, and $T = 1080$ hours.

For the simulation analysis, we test different levels of λ and cost per audit k . To determine parameter values, we consult with a big-box retailer regarding its in-store logistics and interview store managers. Based on operating practices in our research site, we set $\tau = 15$ and examine ($s = 12, S = 24$) under low demand ($\lambda = 1$), ($s = 36, S = 60$) under medium demand ($\lambda = 4$), and ($s = 60, S = 96$) under high demand ($\lambda = 7$). Those parameters are empirically grounded as opposed to analytically optimized due to practical considerations in the field. From store associates we learn that τ is primarily affected by in-store capacity, whereas the target shelf inventory level S is a function of numerous factors. Arguably, if a retailer has excess labor capacity to sustain a much smaller τ and/or abundant shelf space to sustain a much larger S for an item (given the same demand rate), chances are that shelf OOS will be absent and no extra shelf audits (e.g., z^*) will be ever needed. However, in reality many retail stores cannot afford small τ due to understaffing (Chuang et al., 2016), whereas S is limited and constrained by shelf space, merchandising purposes, assortment protocols, and supplier contracts. Thus, first line managers have limited degrees of freedom in reducing τ or augmenting S . Moreover, under limited shelf space, for each item some managers deliberately lower (s, S) in order to enhance product variety and create scarcity effects to stimulate demand. Consequently, the odds of shelf OOS for a SKU increase due to inflexible τ and bounded S . Our audit initiatives driven by zero sales signals exactly respond to the issue and aim to improve on-shelf availability. Following the field work described above, for each run we test four shelf control policies:

- Policy 1 (P1):** (τ, s, S) shelf control.
- Policy 2 (P2):** (τ, s, S) shelf control and $z^*(P(OOS)=w = 0.01)$.
- Policy 3 (P3):** (τ, s, S) shelf control and $z^*(P(OOS)=w = 0.05)$.
- Policy 4 (P4):** (τ, s, S) shelf control and $z^*(P(OOS)=w = 0.10)$.

Policies 2, 3, and 4 differ in their perceived odds of shelf-OOS, and the ones with higher w and $P(OOS)$ (e.g., P3, P4) are more likely to trigger extra shelf audits in the presence of z signals. These three policies add shelf audits driven by z signals onto (τ, s, S) shelf control that already has a periodic counting cycle of τ hours. While those POS data-driven audits would incur extra auditing costs, the reduced lost sales may outweigh the costs. For simplicity we assume $r = 1$ (i.e., no random fixes) for policies 2, 3, and 4.

Fig. 6 shows the average cost per period (averaged over 1000 runs) of policies 1–4 in each of the tested scenarios. Note that the simulation results are qualitatively similar under other realistic parameter settings. As expected, the average cost of all policies increases with cost per audit k . The left panel suggests that under a low demand rate ($\lambda = 1$), the more active P3 and P4 result in slightly higher system costs on average. With that said, the existing (τ, s, S) shelf control policy and the relatively conservative P2 achieve nearly identical cost performance, whereas the latter (P2) helps improve on-shelf availability through triggering extra shelf




Fig. 6. Cost performance of different policies.

audits (z^*). The middle and right upper panels show that when demand rates are medium ($\lambda = 4$) and high ($\lambda = 7$), the extra shelf audits substantially improve cost performance. Even though the most active P4 performs worst in the low demand case, it has the best performance in nearly all of the remaining cases (except the case $k = 1.75$ and $\lambda = 4$). On the other hand, the best-performing P2 under low demand ($\lambda = 1$) is outperformed by the less conservative P3/P4 when demand rates (λ) increase (especially under high demand $\lambda = 7$), indicating the value of data-driven audit initiatives.

The simulation results shown in Fig. 6 suggest that for slow-moving items (with low λ and low shrinkage δ), the extra audits triggered by z signals result in comparable average cost per period. When demand (and hence shrinkage) rates increase, the audits driven by z signals enable retailers to achieve much better cost performance, by correcting shelf failures unfixed within the re-shelving cycle of the (τ, s, S) shelf control. As a matter of fact, the simulation study excludes other potential execution errors (e.g., misplacement, cluttered layout) that harm store operations, and hence the benefits of extra audits could be under-estimated. Nevertheless, the simulation still corroborates the economic value of the proposed shelf audit policy, which can be a cost-efficient complement to existing periodic shelf audit and replenishment operations.

6. Estimation and implementation

After obtaining an understanding of model behaviors (in Sections 3 and 4) and assessing the cost benefits of the proposed model (in Section 5), here we detail two major tasks – estimation and implementation – that are critical to model validation. The first task is aimed for empirically deriving estimates of key model parameters. The second task is focused on examining decisions constructed from the model and efficacy of decisions. Even though the two tasks may not be exhaustive, they have covered the core purposes of our decision support model. Finishing the tasks will be a crucial premise for the POS data analytics approach to be employed in practice.

For the estimation, suppose a store manager has n observations: $(X_1, z_1), \dots, (X_n, z_n)$, where $X_i = 1$ if OOS, $X_i = 0$ if \sim OOS, and z_i denotes consecutive zero sales in POS data, for $i = 1, 2, \dots, n$. The first approach – maximum likelihood estimation (MLE) – is quite straightforward as we can write a joint log-likelihood function $L(w, \hat{p}, P(OOS))$ as

$$\sum_{i=1}^n I(X_i = 1) \log \left(\frac{\widehat{P(OOS)} w \frac{1 - [(1-w)\hat{p}]^{z_i}}{1 - (1-w)\hat{p}}}{\widehat{P(OOS)} w \frac{1 - [(1-w)\hat{p}]^{z_i}}{1 - (1-w)\hat{p}} + (1 - \widehat{P(OOS)}) (1-w)^{z_i} \hat{p}^{z_i}} \right) +$$

$$\sum_{i=1}^n I(X_i = 0) \log \left(\frac{(1 - \widehat{P(OOS)}) (1-w)^{z_i} \hat{p}^{z_i}}{\widehat{P(OOS)} w \frac{1 - [(1-w)\hat{p}]^{z_i}}{1 - (1-w)\hat{p}} + (1 - \widehat{P(OOS)}) (1-w)^{z_i} \hat{p}^{z_i}} \right)$$

where $I(\cdot)$ is an indicator function, and the factor inside $\log(\cdot)$ is $P(OOS|z)$ (from Proposition 1) for $I(X_i = 1)$ and $P(\sim OOS|z)$ for $I(X_i = 0)$. The two estimates $(\widehat{P(OOS)}, \hat{p})$ can be derived before performing MLE. The former $(\widehat{P(OOS)})$ can be estimated by calculating the number of OOS events relative to the number of past audit initiatives. The latter (\hat{p}) can be estimated by computing the ratio of total number of observed zero sales to the number of periods elapsed. The ratio needs to be multiplied by $(1 - \widehat{P(OOS)})$ to recover the intrinsic zero sales probability under \sim OOS. Finally, one just needs to find $\hat{w} = \text{argmax}_{0 \leq w \leq 1} L(w, \hat{p}, \widehat{P(OOS)})$ through numerical maximization of the function L . This MLE protocol is applicable to the case with random fixes of shelf-OOS, where the log-likelihood function is similar to the one above except that $P(OOS|z)$ and $P(\sim OOS|z)$ inside $\log(\cdot)$ are from Proposition 4 (with random fixes). In this case, the goal of numerical optimization is to find $(\hat{w}, \hat{r}) = \text{argmax}_{0 \leq w, r \leq 1} L(w, r, \hat{p}, \widehat{P(OOS)})$.

One may argue that the MLE approach is predicated on pre-derived estimates of $(p, P(OOS))$ that do not fully account for uncertainty in the two parameters. Hence, we propose an alternative Bayesian estimation approach that is more sophisticated yet computationally more intensive. First of all, we define prior distributions for $(w, p, P(OOS))$:

$$w \sim f(\cdot) \\ p \sim g(\cdot) \\ P(OOS) \sim h(\cdot)$$

The three prior distributions can be any parametric distributions that reflect a manager's belief *ex ante*. Since $(w, p, P(OOS))$ range between $(0, 1)$, one could adopt beta or uniform priors. Distributional parameters of $f(\cdot)$, $g(\cdot)$, and $h(\cdot)$ depend on managers' subjective beliefs and previous observations. Given the priors, we can derive full conditional distributions known to a certain proportionality for $(w, p, P(OOS))$:

$$P(w|\mathbf{X}, \mathbf{z}, p, P(OOS)) \propto f(\cdot) \prod_{i=1}^n P(X_i|z_i, w, p, P(OOS)) \\ P(p|\mathbf{X}, \mathbf{z}, w, P(OOS)) \propto g(\cdot) \prod_{i=1}^n P(X_i|z_i, w, p, P(OOS)) \tag{4} \\ P(P(OOS)|\mathbf{X}, \mathbf{z}, w, p) \propto h(\cdot) \prod_{i=1}^n P(X_i|z_i, w, p, P(OOS))$$

where \mathbf{X} and \mathbf{z} refer to the vectors of observations $(X_1, z_1), \dots, (X_n, z_n)$. $P(X_i|z_i, w, p, P(OOS))$ is $P(OOS|z)$ and $P(\sim OOS|z)$ (from

Proposition 1) for the i th pair of observations. Given the three full conditional distributions in (4), we can apply the Metropolis-Hastings-within-Gibbs sampler (Hoff, 2009) – a popular Markov chain Monte-Carlo (MCMC) algorithm – to construct the posterior distributions. Take w for instance, the sampled $w^{(s)}$ transits to $w^{(s+1)}$ in the following way:

1. Define a proposal distribution $J(w|w^{(s)})$
2. Sample a proposal value w^* from $J(w|w^{(s)})$
3. Compute the acceptance ratio $\phi = \frac{J(w^{(s)}|w^*)P(w^*|\mathbf{X}, \mathbf{z}, p, P(OOS))}{J(w^*|w^{(s)})P(w^{(s)}|\mathbf{X}, \mathbf{z}, p, P(OOS))}$
4. Let $w^{(s+1)} = \begin{cases} w^* & \text{with probability } \min(\phi, 1) \\ w^{(s)} & \text{with probability } 1 - \min(\phi, 1) \end{cases}$

For the proposal distribution $J(\cdot)$, one can employ a random walk proposal (Brooks, 1998) to initialize the Bayesian simulation. The acceptance ratio ϕ is based on the prior $f(\cdot)$ and the full conditional of w . The Monte-Carlo sampling for p and $P(OOS)$ are done with their respective priors and full conditionals in a similar fashion. After S iterations, the sequences $(\mathbf{w}, \mathbf{p}, \mathbf{P}(OOS))$ are expected to form a stationary Markov Chain and constitute the posterior distributions that we look for. The simulation convergence (desired Markovian behaviors) – irreducible, aperiodic, and recurrent – can be evaluated by standard metrics for MCMC (e.g., stationarity and no stickiness). The posterior means/modes of $(\mathbf{w}, \mathbf{p}, \mathbf{P}(OOS))$ serve as robust estimates of $(w, p, P(OOS))$ to which optimal decisions are sensitive. The parameter r (in the case with random fixes) can also be estimated by the Bayesian methodology with a little modification (i.e., an addition prior for r and $P(OOS|z)$ from Proposition 4 as opposed to Proposition 1). Moreover, the Bayesian method above is particularly valuable to practical implementation of the data-driven audit policy, since in reality the two Markov chains in Sections 3.1 and 4.1 could be non-stationary. That is, parameters w and r in the two transition matrices may change over time. The nature of Bayesian update – constantly refine prior beliefs and posterior estimates using latest data – allows decision-makers to continuously update their estimates of both $(p, P(OOS))$ and (w, r) that govern Markov chains.

After empirically deriving estimates of probabilistic parameters $(w, r, p, P(OOS))$ (from either MLE or Bayesian estimation), cost factors (k, m) , and zero-truncated mean μ^+ (that can be estimated by computing sample mean of non-zero POS observations), the proposed shelf audit policy can be implemented in actual operations. To begin with, a decision-maker needs to select one out of the two decision support models (i.e., with or without random fixes). To do so, he/she can use the estimated values of $L(\hat{w}, \hat{p}, \hat{P}(OOS))$ and $L(\hat{w}, \hat{r}, \hat{p}, \hat{P}(OOS))$ to derive a likelihood ratio test (a χ^2 test with $df = 1$) for $H_0: r = 1$ (no random fixes). If no significant evidence is found to reject $H_0: r = 1$, it will be safe to go for the first model in Section 3.

The last element of implementation is to validate effectiveness of actual audit decisions. Based on the outcomes of n audit initiatives, managers can quickly learn how many audits fix shelf-OOS (b) and how many audits turn out to be false alarms ($n-b$). A high b/n indicates that the selected model (with or without r) is valid and effective. On the other hand, a low b/n would be a strong indicator of overly sensitive audit triggers. The assessment of b/n will give managers a clear idea of economic benefits (in terms of the number of fixed stock-outs) and operational feasibility. In either high or low b/n , a manager is supposed to use observed audit outcomes to regularly re-estimate model parameters $(w, r, p, P(OOS))$ (using the estimation method shown earlier), in addition to constantly examining decisions constructed from the model and efficacy of decisions, such that the model will be valid in its practical use.

7. Concluding remarks

All parties in a retail supply chain need to realize that shelf-OOS and the consequent low product availability undermines not only operational but also financial performance. The prevalence of OOS forces retailers to occasionally perform manual audits to enhance product availability. Even though RFID-enabled automatic counting is an attractive alternative to manual audits (Zhou, 2009), a complete deployment of item-level RFID is still hard to achieve for many retailers. When done properly, the old-fashioned shelf audits can be *efficacious* (reducing shelf-OOS), *efficient* (done economically), and *effective* (elevating sales) (Chuang et al., 2016). However, shelf audit decisions need to be made in an evidence-based and cost-informed fashion. Store managers usually look for OOS evidence from POS records but do not have a formal optimal criterion for triggering audit initiatives. In response to managers' need for improving shelf audit decisions, our paper presents a POS-data analytics approach for managers who aim to fix the recurring issue of shelf-OOS through manual audits.

Our data-driven audit policy only uses consecutive zero sales observed directly in POS data as the state variable and imposes no peculiar assumptions on replenishment policies. We avoid using system inventory level as a state variable because inventory records in retailing are largely inaccurate (Chuang & Oliva, 2015; DeHoratius & Raman, 2008). The model is simple and practically easy in the sense that it just takes relevant cost factors, conditional probability of shelf-OOS, and conditional expectation of shelf-OOS duration as inputs. That said, the model is also complex and theoretically sophisticated as it considers intrinsic sales potential, zero sale probability of demand, stochastic transition from non-OOS to OOS, perceived OOS likelihood, and even random fixes of OOS. In labor-intensive retail operations, store employees could fix some OOS when walking through the aisles and seeing empty shelves. Despite increasing model complexity, the incorporation of random fixes into shelf audit policy design distinguishes our work from prior studies and makes our modeling effort better reflect reality.

The proposed likelihood framework in Section 6 also greatly enhances the empirical base of our decision model. While analytical modeling approaches ensure cost optimality, parameters of analytical models may not be entirely observable or estimable, and hence heuristic approaches are seemingly more useful for supporting shelf audit decisions (Papakiriakopoulos et al., 2009). We address this issue by showing how to apply the conditional probability to classical MLE and derive parameter estimates. Further, we demonstrate the use of Bayesian and MCMC methods to tackle uncertainty in unobservable model parameters. Being able to estimate stochastic failures (w) and random fixes (r), both of which have potential non-stationarities, is critical to decision quality because ignoring those factors would bias $P(OOS|z)$ estimates as well as decisions resulting from the model. Bayesian inference has been adopted by operations researchers to infer demand parameters (Hill, 1997) and inventory level (DeHoratius et al., 2008) in order to improve replenishment decisions. However, Bayesian hierarchical model is rarely applied due to its computational complexities, which are less of an issue nowadays due to recent advances in sequential Monte-Carlo methods. Our paper shows a compelling example of using Bayesian methods to develop a statistically grounded decision support model. This venue is promising as the complementarity between statistical computing and optimization modeling would enable operations researchers to better leverage data in models.

In an attempt to validate the cost-efficacy of monitoring shelf-OOS and activating audits based on z signals, we conduct simulation experiments in addition to numerical studies on model behaviors. Simulation results show that decoupling shelf audit policy design from replenishment reduces model complexity with-

out harming policy applicability. Even though there are instances where the data-driven audits would result in slightly higher average costs, in most of the tested cases the proposed model would substantially improve cost performance. The simulation study verifies that our modeling effort based on irregular anomalies (i.e., consecutive zero sales) can be used to run alongside retailers' regular operations (i.e., periodic shelf audit and replenishment) in a cost-effective fashion, especially for non-slow moving items. In a nutshell, the data-driven model for shelf audit decisions is easy to understand, well parameterized, and consistent with the requirements of a good 'decision calculus' (Little, 2004). Being independent of any inventory replenishment policies (e.g., EOQ, base-stock), our POS data analytics approach potentially can be applied to most retailing contexts where shelf-OOS is an outstanding issue.

References

- Astrup, J., & Kotzab, H. (2010). Forty years of out-of-stock research – And shelves are still empty. *The International Review of Retail, Distribution and Consumer Research*, 20(1), 147–164.
- Brooks, S. P. (1998). Markov Chain Monte Carlo method and its application. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 47(1), 69–100.
- Cachon, G. (2012). What is interesting in operations management. *Manufacturing & Service Operations Management*, 14(2), 166–169.
- Chuang, H. H. (2015). Mathematical modeling and Bayesian estimation for error-prone shelf audits. *Decision Support Systems*, 80, 72–82.
- Chuang, H. H., & Oliva, R. (2015). Inventory record inaccuracy: Causes and labor effects. *Journal of Operations Management*, 39–40, 63–78.
- Chuang, H. H., Oliva, R., & Liu, S. (2016). On-shelf availability, retail performance, and external audits: A field experiment. *Production and Operations Management*, 25(5), 935–951.
- Condea, C., Thiesse, F., & Fleisch, E. (2012). RFID-enabled shelf replenishment with backroom monitoring in retail stores. *Decision Support Systems*, 52, 839–849.
- DeHoratius, N., & Raman, A. (2008). Inventory record inaccuracy: An empirical analysis. *Management Science*, 54(4), 627–641.
- DeHoratius, N., Mersereau, A. J., & Schrage, L. (2008). Retail inventory management when records are inaccurate. *Manufacturing & Service Operations Management*, 10(2), 257–277.
- Dudley, R. (2014). Wal-mart sees \$3 billion opportunity refilling empty shelves. Bloomberg News. <http://www.bloomberg.com/news/2014-03-28/wal-mart-says-refilling-empty-shelves-is-3-billion-opportunity.html>.
- Fisher, M. L., & Raman, A. (2010). *The new science of retailing: How analytics are transforming the supply chain and improving performance*. Boston: Harvard Business Press.
- Greenhouse, S., & Tabuchi, H. (2014). Wal-mart memo orders stores to improve grocery performance. The New York Times. http://www.nytimes.com/2014/11/12/business/walmart-memo-orders-stores-to-improve-grocery-performance.html?_r=0.
- Gruen, T. W., & Corsten, D. (2008). *A comprehensive guide to retail out-of-stock reduction in the fast-moving consumer goods industry*. Arlington, VA: Food Marketing Institute.
- Hill, R. M. (1997). Applying Bayesian methodology with a uniform prior to the single period inventory model. *European Journal of Operational Research*, 98(3), 555–562.
- Hoff, P. D. (2009). *A First Course in Bayesian Statistical Methods*. New York: Springer.
- Iglehart, D. L., & Morey, R. C. (1972). Inventory systems with imperfect asset information. *Management Science*, 18(8), 388–394.
- Kang, Y., & Gershwin, S. B. (2005). Information inaccuracy in inventory systems: Stock loss and stockout. *IIE Transactions*, 37(9), 843–859.
- Kumar, S., & Arora, S. (1992). Development of internal audit and cycle-counting procedures for reducing inventory miscounts. *International Journal of Operations & Production Management*, 12(3), 61–70.
- Kok, A. G., & Shang, K. H. (2007). Inspection and replenishment policies for systems with inventory record inaccuracy. *Manufacturing & Service Operations Management*, 9(2), 185–205.
- Little, J. D. C. (2004). Comments on "models and managers: The concept of a decision calculus. *Management Science*, 50(12), 1854–1860.
- Nachtmann, H., Waller, M. A., & Rieseke, D. W. (2010). The impact of point of sale data inaccuracy and inventory record data errors. *Journal of Business Logistics*, 31(1), 149–158.
- O'Reagan, R. T. (1969). Relative costs of computerized error inspection plans. *Journal of the American Statistical Association*, 64(328), 1245–1255.
- Papakiriakopoulos, D., Pramatari, K., & Doukidis, G. G. (2009). A decision support system for detecting products missing from the shelf based heuristic rules. *Decision Support Systems*, 46(3), 685–694.
- Sandoh, H., & Shimamoto, H. (2001). A theoretical study on optimal inventory-taking frequency for retailing. *Journal of Retailing and Consumer Services*, 8(1), 47–52.
- Ton, Z., & Raman, A. (2010). The effect of product variety and inventory levels on retail sales: A longitudinal study. *Production and Operations Management*, 19(5), 546–560.
- Zhou, W. (2009). RFID and item-level information visibility. *European Journal of Operational Research*, 198(1), 252–258.
- Zondag, M. M., & Ferrin, B. (2014). Finding the true voice of the customer in CPG supply chains: Shopper-centric supply chain management. *Journal of Business Logistics*, 35(3), 268–274.