Point-of-use hospital inventory management with inaccurate usage capture

Benjamin V. Neve[1] · Charles P. Schmidt[2]

Abstract

Many hospital supply chains in the US follow a "stockless" structure, often implemented with the acquisition of new systems promising improved efficiencies and responsiveness. Despite vendor promises, supply chain gains from new technology are often unfulfilled or result in a reduction of performance. A critical component of achieving promised gains is the hospital's ability to accurately and consistently capture hospital inventory use. In practice, recording demand with perfect, 100% accuracy is infeasible, so our models condition on the level of accuracy in a particular hospital department, or point-of-use (POU) inventory location. Similar to previous literature, we consider *actual* net inventory and *recorded* net inventory in developing the system performance measures. We develop two models, optimizing either cost or service level, and we assume a periodic-review, base-stock (or par-level) inventory policy with full backordering. In addition to choosing the optimal order-up-to level, we seek the optimal frequency of inventory counts to reconcile inaccurate records. Results from both models provide insights for supply chain managers in the hospital setting, as well as hospital administrators considering the adoption of similar technologies or systems.

Keywords Hospital supply chain · Inventory management · Inaccuracy · Internal logistics · Operations research

Article Highlights

- As healthcare decision-makers consider inventory levels and their related inventory systems, inaccuracy at the point-of-use can create significant operational hurdles when trying to either reduce costs, or to maintain high service levels.
- We present two inventory decision models that consider inaccuracy, specifically in a hospital setting where front-line patient care staff use an open-bin barcode or button-scan inventory system prevalent in hospitals in the United States, which is also tied to patient billing. These specific types of inventory systems have not been studied through the lens of inaccurate records elsewhere in the literature, though inaccuracy has been considered in other models.

- The models presented both perform well in seeking to optimize costs and meet service-level requirements, showing that it is possible to still provide timely patient care without perfectly accurate inventory – though the costs and operational impacts can be significant as increased inventory levels will use up more valuable hospital space, and increased counting frequencies will have a direct impact on labor costs.
- While the structure of the models are not unimodal, the incremental solution process is shown to perform well, finding the optimal solution 98% of the time in the shortage cost model and 74% of the time in the service level model. Non-optimal solutions were less than 2.4% away on average from optimal values, with the shortage cost model showing better performance.
- This research should give insight to both hospital inventory managers and medical clinicians as they consider their joint role in ensuring patients have access to the supplies necessary for timely and effective care. Additionally, operational symptoms described within the paper may direct decision-makers to consider the

✉ Benjamin V. Neve
benjaminneve@weber.edu

Charles P. Schmidt
cschmidt@cba.ua.edu

[1] Weber State University, 3848 Harrison Blvd, Ogden, UT 84408, USA

[2] Professor of Operations Management, Emeritus, Tuscaloosa, AL, USA

potential impact that inaccuracy is causing on their hospital supply chain and how to address the impact.

1 Introduction

Inaccurate inventory records can lead to increased costs, unexpected delays in demand fulfillment, and lost revenue. Inventory record inaccuracy is also one of the many sources of unnecessary health care costs across the US, and creates the additional risk of delayed patient care. This research presents a model based on personal experiences by the authors with the inventory system of a 580-bed Alabama hospital, where inventory record inaccuracy was found to significantly affect the supply chain, leading to both overstock and daily shortages throughout the facility. This situation led the Materials Manager to adopt a daily inventory counting policy requiring additional employees, yet still failed to resolve the inventory problems. In this case, inaccurate records failed to account for nearly half a million US dollars per year in consumed medical supplies, which represents a patient reimbursement value of about two million US dollars annually. This financial impact is especially risky, as some hospitals experience a larger impact on the cash conversion cycle from inventory issues rather than patient care issues, especially during the COVID-19 pandemic [32].

Many of the hospital's staff associated with patient care do not feel responsible for inventory management, and yet expect inventory to always be available when needed - a situation which is a common reality in US hospitals [25]. Conversely, the material management staff are tasked with (and measured against) ensuring inventory is available at each point of care, so that patients receive care in a manner and time that is most beneficial. Their ability to optimize inventory availability hinges on the accuracy of recorded inventory use as consumed by patient care.

Our research here intends to provide two structured models that can give insight to healthcare decision-makers seeking to maintain service levels or to minimize costs when inventory record accuracy is causing logistical problems. Our models attempt to capture, and demonstrate, the necessity for patient care staff to follow inventory management standards - specifically, to accurately record inventory usage by properly using the system(s) installed in each care location. As 100% accuracy is not always feasible, our models successfully balance cost and service levels through inventory policy decisions in light of patient care staff compliance levels.

The situation we experienced in Alabama was the main motivation behind our research, however, our experience suggests this situation among hospitals is not unique, as bar-code or button-scan inventory systems are used in hospitals across the United States and are heavily reliant on patient care staff to accurately capture inventory use. Research is being done now to study how wide-spread these issues are occurring, as well as the impact on costs and stock availability within healthcare systems more broadly.

1.1 The hospital supply chain

Traditionally, hospitals managed two tiers of their supply chain, where one was internal and the other external [24]. The internal supply channel includes inventory and logistics from the hospital's on-site warehouse to the several point-of-use (POU) inventory locations within the hospital. The external supply channel includes the flow of inventory from the supplier to the hospital's on-site warehouse as shown in Figure 1. In such cases, inventory control policies are required both for the warehouse and for each of the POU locations.

Rivard-Royer et al. [24] discuss that, during the 1990s, hospitals in the US and Canada moved away from the classic hospital supply chain approach, and trended toward a stockless management system, as portrayed in Figure 2.

The stockless system transforms the external supply channel into a direct replenishment stream from the supplier to the POU inventory locations, essentially removing the need to keep stock in an on-site warehouse. Based on direct demand information from each POU, the supplier ships materials in unit quantities directly to the POU. Stockless systems have been shown to drastically reduce inventories by as much as 70% to 80%, as well as reduce the material handler full time equivalents (FTEs) by as much as 30% to 45% [24]. This accounts for several hundreds of thousands of dollars in annual savings for the hospitals involved.

However, the stockless system is not without its drawbacks. It requires up-to-date and accurate transmission of supply needs, a short lead time to ensure minimal inventory at each POU (or sufficient inventory capacity at each POU), and an ability to adapt to variation in product mixes and seasonality of demand. In many instances, hospitals are required to make costly systems upgrades or expand POU inventories to qualify for the stockless system. In other cases, hospitals that experience relatively longer lead times from the supplier are required to keep large amounts of inventory on-hand between available shipments. Rural hospitals and not-for-profit hospitals fall into the group of hospitals that may have difficulty in implementing or affording the changeover to a stockless system. [24]

Other researchers have approached the topic of inventory management in hospitals, with a similar structure studied in [15], where they compare costs between automated and traditional POU inventory systems. More generally, [35] and [18] address issues in the hospital supply chain from an empirical standpoint by surveying representatives from

Fig. 1 Classic hospital supply chain
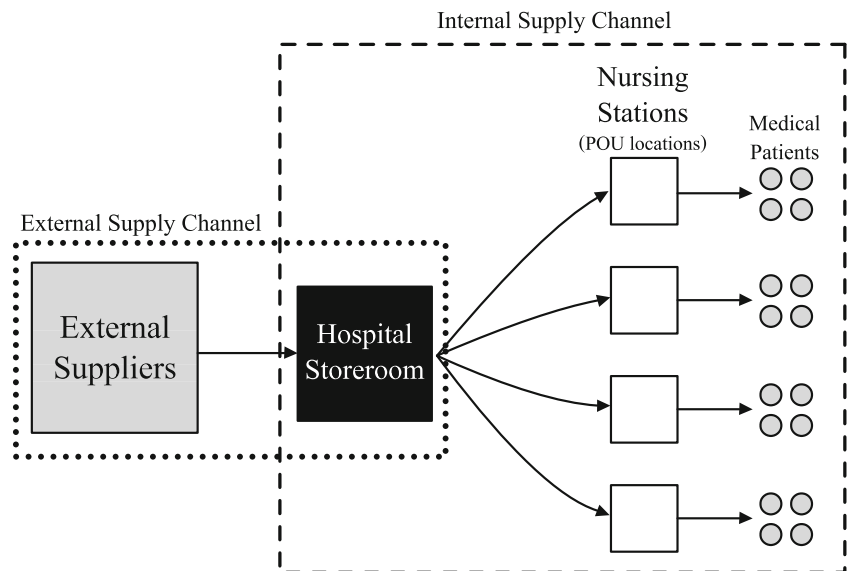
several hospitals. Both empirical studies cite a need for additional research regarding the hospital supply chain, as well as the complexity of such supply chain systems.

Some researchers, such as [8, 9, 11, 13, 15, 26] and [12] provide specific case studies or numerical analyses, where a hospital, or hospitals, are measured before and after an implementation of different supply chain modifications. While insightful, researchers fail to address inaccuracy or optimization of the system. For example, [15] considers a supply chain structure almost identical to the one considered here, however, instead of formulating a model and seeking optima as is done here, they compare explicit costs and recommend using an Economic Order Quantity (EOQ) approach to select control parameters. The simulation

model provided by [11] focuses on medication inventory management and compares the results across four different inventory policies that use varying weights on criticality, availability and expiration window to determine optimal policies. While information accuracy is mentioned in the paper, the authors refer to pharmacy order information rather than accuracy of available on-hand stock. A more recent article from this group focuses on the application of Lean Thinking to improve inventory control in a specific case, but also suggests the improvements only work with adequate information sharing without a solution to address information accuracy [12].

An important finding is presented by [13], where the human "intuition" approach does not perform near as

Fig. 2 "Stockless" hospital supply chain
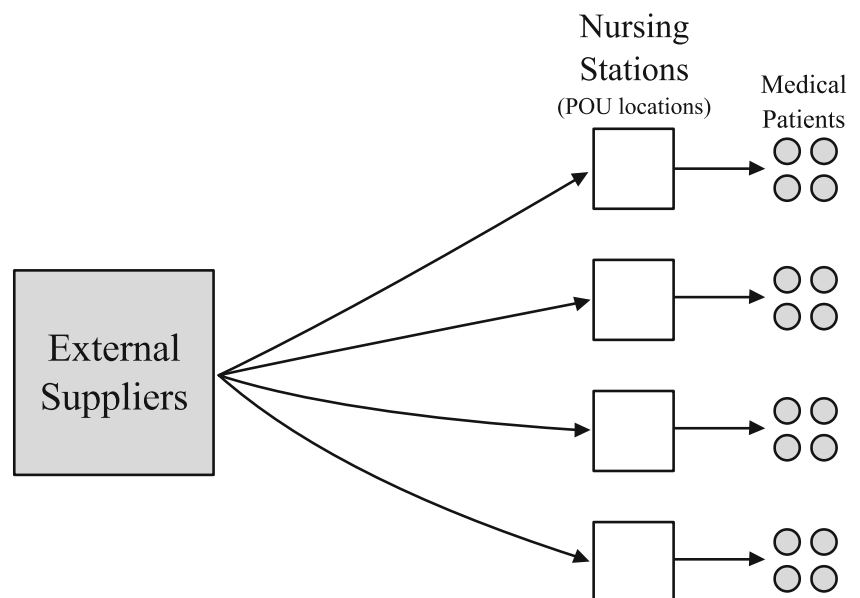
well as optimized decisions in a hospital setting - with their models showing that optimal policies generated by expert systems improve the outcomes by as much as 47%. However, this level of improvement can only be realized if decision makers (or their expert systems) have visibility into actual availability and usage patterns. Similar to our experiences, the author suggests that even with an expert system that can provide better inventory policies, many in the hospital setting do not trust the system and instead circumvent the suggested expert policies, leading to a decline in service levels and increased costs. In practice, inaccurate inventory led to a distrust of the system and to disregard for expert system suggestions. Our models provide a structure and an example of how inaccuracy impacts performance, but also how awareness of the issue can inform decision makers - or designers of expert systems - to improve accuracy, counting policies and/or inventory replenishment policies.

Recent papers that have focused on the optimization of hospital inventory include [1, 3, 4, 6] and [27]. A Markov decision approach is used in [3] to optimize two models, one focused on meeting specific service level requirements, and the other maximizing service levels within constrained inventory storage space. The model presented in [1] considers storage capacity and physical inventory size for a group of stockless items. The model uses average daily inventory usage for each item rather than stochastic demand, as is used here. Similarly, a simulation optimization model is presented by [4], but their focus is on highly specialize inventory related mainly to high-risk heart surgery. The recent work by [6] looked at how the number of Kanban bins affects the optimal inventory solution using a mixed-integer model and deterministic demand. In this recent research, the models explicitly rely on the assumption that the behavior of hospital staff maintains perfect operational conditions - in that the authors hope that the nurses do not misplace Kanban cards. The authors suggest a future opportunity to research how staff behavior may affect hospital inventory models. Lastly, the authors in [27] present and compare robust stochastic inventory models in an Automated Dispensing Machine (ADM) environment within a hospital, where different replenishment mechanisms are considered. The main focus of their paper is the concept of dual-sourcing, and the authors find that coordinated joint replenishment policies perform better than hybrid policies that are not coordinated. In addition to the models presented, [27] discuss how increased inventory visibility can improve inventory outcomes, though this is tied to assumptions about ADM machines, which in some cases, are cost-prohibitive in practice. Overall, though the importance of hospital inventory is explored in each of these models, none consider the direct impact of inaccuracy as a part of their inventory models the health care setting.

For comprehensive reviews of health care related inventory or supply chain models, see [2, 19, 28, 29, 34] and [12]. In [34], the authors focus on literature that is quantitatively oriented and suggest several areas of potential future research. Research focusing on POU inventory models with unpredictable demand is a future research opportunity, and our models here provide an example of this. The authors in [34] also suggest that there is a lack of research addressing up-to-date stock information. Looking at the review provided by [19], there is potential for new research to develop more effective internal supply chain performance measures that lead to overall improvement of internal hospital logistics. The authors also indicate that there remains difficulty within industry of maintaining up-to-date inventory management parameters. Beyond the case study discussed previously, [12] also provides a comprehensive, well-organized list of articles that address healthcare inventory models in recent years.

Additional, the review provided by [29] captures a broad perspective of analytical modeling techniques as they are applied to hospital inventory systems, considering both costs and service-levels. Their paper includes a classification system that provides an overview of modeling techniques, solution approaches and a variety of applications in a hospital setting. Our paper would be classified as a stochastic programming model, where both the demand and the accuracy of recording demand are represented with stationary stochastic distributions.

The authors in [2] focus their review on papers that study surgical supplies and sterile instruments destined for operating rooms, with an emphasis on models that seek cost savings. They suggest that previous articles can be classified as either planning and scheduling or as inventory management research. Hospital inventory models that account for operational or disruption risk were identified as an important future research opportunity. A similar research opportunity was suggested by [29], where research that considers behavioral uncertainties is lacking.

To explore the future opportunities suggested by [2] and [29], our research presents two inventory management models to address the situation when inventory-tracking behavior of medical staff can lead to disruption or operational risk.

Additionally, we find that none of the reviews mentioned research that has been done to address the inaccuracy of inventory within the hospital setting, so our research seeks to fill this gap. In the next section, however, we will introduce some potentially overlooked literature focused on inventory inaccuracy, including a doctoral thesis that considers inaccurate inventory in a healthcare setting [21].

1.2 Inaccurate inventory

Although there are not many hospital-centric inventory models that address inaccuracy, the generalized concept is not new to the literature [14]. We categorize related generalized literature into two main categories: analysis or optimization.

Analysis-focused literature is either empirical in nature, or it does not present a specific model nor its optimization. Empirical research focusing on inaccuracy can be found in [20, 31] and [22], while [10] present a simulation study and share results and insights. One article found that inaccuracy amounted to 10% of profits over a certain period for a select case [22].

The simulation developed by [10] studied the effect of inventory inaccuracy on the performance of a three echelon supply chain. In their model, they consider several factors that may cause inventory inaccuracy, but do not include transaction errors. Similarly, [20] focus on the situation where recorded inventory is greater than actual inventory (overstatement error). They describe how overstatement errors can cause a system to postpone replenishment, leading to lower service levels. While our service level model provides similar insights to [20], we also provide a cost model that considers the impact of inventory inaccuracy on overall costs and the resultant service levels. Additionally, our research assumes a discrete Poisson demand, rather than a continuous normal distribution as used previously in [20]. Our insights also provide more targeted guidance to the healthcare field, as our models are informed by the unique situations within hospital supply chains.

Optimization-oriented literature addressing inaccuracy includes work spanning 1972 through 2020. We have included references to three types of optimization models: replenishment policy-focused [23], inspection policy-focused [30], and those that include both types of policy decisions [7, 14, 16, 21] and [5].

A newsvendor model developed by [23] focused on a retail store where they optimize the replenishment policy when inaccurate inventory is considered in one case, and where the firm is assumed to have perfect inventory information in another case, due to utilizing a Radio Frequency Identification(RFID) tracking system.

Conversely, [30] ignores replenishment decisions to focus on the costs of both counting inventory and investigating the causes of inventory discrepancies in a supermarket or convenience store environment. They differentiate the objective to find the minimum cost, giving the optimal cycle count frequency between annual accounting audits.

Lastly, we discuss articles that present inventory models where replenishment and inspection policies are adjusted to address inaccurate records, as it represents research that is most similar to our own.

The first quantitative research regarding inaccurate inventory records was provided by [14], where they develop an optimization model with the assumption that a replenishment policy is given. However, they modify the given replenishment policy by setting a buffer stock level sufficient to cover all probable inventory record errors until the next inventory count. In their model, [14] do not allow for backordering and they assume the magnitude of record errors is not dependent on the magnitude of demand. In our paper, errors are handled using a probability that particular demand event is recorded, though when physical stock is not available all demand is backordered. As such, we assume full backordering, and we seek an optimal base stock level rather than a buffer stock quantity under a given replenishment policy.

Similar to our own research, [16] develop a model for inaccurate inventory where they make joint inspection and replenishment decisions to minimize system costs. The inspection decision is triggered dynamically, and it perfectly aligns the inventory record with on-hand inventory. A base-stock replenishment policy is shown to be optimal, where both replenishment and inspection are tied to the level of inaccuracy in the system. By comparison, the research we present here also assumes a base stock replenishment policy, but we choose a static counting frequency before the planning period begins, which is different than in [16] who utilize a dynamic inspection rule. Furthermore, our research includes an additional model where we seek to optimize the inventory system based on service-level requirements rather than on total cost.

Similar to a previously cited paper by [23], the researchers in [7] develop a model based in a retail setting, but where they propose a modified inventory record using a Bayesian approach. Their research is different from our own, as they assume inventory inaccuracy with lost sales (instead of backordering) as they attempt to optimize the auditing frequency and replenishment policy.

Research conducted in the North East United States provides a comprehensive, multi-method approach to study inaccuracies in a hospital inventory system utilizing Automated Dispensing Machines (ADMs) [21]. The author includes an empirical study to classify the sources of inventory inaccuracy, a simulation that models the inventory system to understand the effect of the inaccuracies on product availability, an analytical model that demonstrates a generalization of the simulation study, and an applied case study for implementing improvements at a specific hospital.

While we do consider inaccuracy as a significant influence on product availability, we do not consider ADMs in our model. Thus, our research provides insight for hospitals that use a bar code or button scanning system in an open-inventory environment - as is often used for standard inpatient care inventory stores at the POU. Furthermore, we present a model that also considers the cost implications of inaccurate records, in addition to a model that considers inventory availability as in [21]. Aside from our research, [21] is the only research we could find that considered inaccurate inventory in a hospital setting, but they do not consider costs as we do, and their research depends on data from an ADM system whereas our model better mimics the prevalent barcode or button-scan systems with open bins at the POU.

Most recently, [5] provides another example of a study that seeks to understand the impacts of inventory inaccuracy in a general setting. The authors assume an (R, Q) policy with lost sales, and allow for *incoming* and *outgoing* transaction errors, random demand and uncertain supply and lead time. They provide two models, one with the goal of adjusting safety stock to counter inaccurate inventory, and the other to adjust the recorded inventory level in an attempt to trick the system to maintain sufficient on-hand inventory. Their model also includes an analysis of a triggering process to initiate inventory counts for reconciliation of records (similar to [16]). They use a best case and a worst case to give boundaries to a detailed analysis of how inaccuracy affects the supply chain, suggesting that improved methods of addressing inventory inaccuracy can decrease total inventory levels while maintaining sufficient stock to satisfy demand. They suggest the need for a cost-based structure for future research, which our model does consider.

2 General model assumptions

After reviewing the literature, we have found our research to be novel. We are the first to propose an inventory model for a hospital supply chain that addresses inaccuracy, when inventory is managed by a bar code or button-scan electronic inventory system at the POU - which is a widely used system in the health care industry. (The open-bin electronic tracking system considered here is different than ADM systems that securely manage access to inventory within an automated locking cabinet.) Among the inaccuracy-related inventory literature, we are the first to seek improvements to both replenishment and inspection policies where 1) the random demand is assumed to be discrete, 2) transaction errors are assumed to generate only positive (overstatement) error in recorded inventory, 3) both a cost-minimization model and

a service-level model are presented, and 4) the situation is based on the health care industry.

We assume the hospital operates under a stockless supply chain structure, as given by Figure 2. As is the case in the Alabama hospital, nurses and staff at the POU inventory locations are responsible for recording supply usage through the use of an electronic inventory tracking system that utilizes button-scan technology. If the user fails to properly record the consumption of a particular supply item, the *recorded* inventory level within the electronic system remains unchanged while *actual* available inventory decreases. In the literature, this type of error has been called an overstatement error, where the recorded system inventory level is consistently greater than or equal to the actual inventory level.

In an overstatement-error situation, the system attempts to order up to the par level for the next period, but the *recorded* inventory level shows more inventory on-hand than is actually on hand. For example, an inventory item with a par level of 30 might show 25 units on hand in the system, but may only have 12 units in the bin because 13 units were consumed without being recorded. The system would only order 5 units, instead of 18 units - leading to a likely stockout and backorder event(s) during the next period.

We define p to represent the probability that a demand event is successfully recorded, and we call this probability the *accuracy level* or the level of recording accuracy. For example, if the accuracy level is 65% (i.e. $p = .65$), then there is a 65% probability that demand is recorded. This means that the recorded net inventory will be greater than actual inventory, since actual inventory will decrease at a faster rate than recorded inventory.

In an attempt to create a stockless hospital supply chain, a hospital may desire to use the electronic inventory system mentioned above to automatically place regular orders for replenishment, as well as to simultaneously assign inventory items to patient bills - assuming patient care staff successfully record use. However, when the system is inaccurate, orders are delayed and the quantities ordered may be insufficient, causing backorders. It is important to note that it is not the backorders that equate to failed reimbursement, but rather, when an item used in patient care is not recorded, and thus, not assigned to a patient bill. In other words, the patients cannot be charged for the use of that supply item, as there is no record generated on their patient bill. This is how inaccurate inventory records lead to failed reimbursement, to ineffective ordering, to stockouts, and sometimes to delays in patient care.

We now present two inventory models for a hospital that is attempting to operate according to a *stockless* supply chain, but must deal with inaccurate inventory

records (given an accuracy level, p). We assume a periodic-review base stock inventory replenishment policy with full backordering, and we seek the optimal base stock level to minimize costs in the first model. In the second model, we seek the smallest base stock level required to meet a predetermined customer service level.

In Section 3 we present the cost-minimization model with its optimal solution. In Section 4 we discuss the service-level model with a solution approach. In Section 5 we discuss the solution approaches for both models, including heuristics for near-optimal solutions. Section 6 we present the results of seeking optimal and near-optimal solutions for both models. We conclude in Section 7 with an Appendix in Section 8.

3 Cost model

In this section, we introduce the inventory model with a cost objective which we will later seek to minimize. First, we discuss the physics and structure of the supply chain while introducing some of the notation and assumptions. Then we present the formulation with the cost objective we later minimize.

3.1 Model structure

As stated previously, the following model is based on the inventory operations of the Alabama hospital that motivated this work. We develop the periodic-review inventory model for a single item at a single POU location, where replenishment comes directly from an external supplier. Thus, the replenishment of POU inventory bypasses the hospital's main storeroom.

To model the described inventory system, we assume that time is divided into daily periods, where each period, t, represents a 24-hour time interval. Orders are assumed to arrive 24 hours after being placed, so lead time, L, is one period, $L = 1$. At the beginning of each period, the POU receives the incoming order from the supplier. Immediately following order receipt, an order is placed with the external supplier through an automated system. Therefore, both receiving and ordering occur at the beginning of every period and in that sequence, respectively. This is basically what occurs in practice, but the events do not occur instantaneously, as is assumed here.

The automated replenishment system is responsible for electronically-generating orders each period with the external supplier according to a base stock policy with order-up-to level S (often referred to as the *par level*). The actual order quantity is based on the difference between the current *recorded* inventory level, I', and S. For explanatory

purposes, we define $\bar{I}'_i$ to be the recorded inventory at the end of period i. We also define $\underline{I}'_i$ to be the recorded inventory at the beginning of period i, *after* the outstanding order arrives. That means that when the system places an order at the beginning of period i, it will place an order of size $S - \underline{I}'_i$. This order will arrive at the beginning of the following period.

During a particular period(day), nurses and nursing staff personally remove inventory from the POU to use for patient care. Upon removing the inventory, the person taking the supplies is supposed to manually record which item is taken, and which patient will be billed for the item. There are several different automated POU inventory systems on the market, each with its own method for recording usage - some easier than others. At this hospital, the recording process takes a matter of seconds, but it requires the use of a "data wand" that is used to touch preprogrammed e-tags on the patient charge board and on the inventory supply bins. We have already defined p to be the probability that this usage-recording process occurs for a particular demand instance.

So, during each period, i, there is a portion of demand that is successfully recorded, which we define as the *recorded* demand in period i, RD_i. It is crucial to note, that the recorded inventory, I', is depleted only by recorded demand. This gives the following relationships, which are given for explanatory purposes:

$$\bar{I}'_i = \underline{I}'_i - RD_i \tag{1}$$

$$\underline{I}'_{i-1} = S - RD_{i-2} \tag{2}$$

$$\underline{I}'_i = \bar{I}'_{i-1} + (S - \underline{I}'_{i-1}) = \bar{I}'_{i-1} + RD_{i-2} \tag{3}$$

Equation 1 shows how recorded inventory at the end of period i is affected by recorded demand. It is simply the recorded inventory at the beginning of period i, less the demand recorded during period i.

Equation 2 gives the relationship of the previous period's beginning recorded inventory with the prior period's recorded demand. The recorded demand quantity for period $i - 3$ was ordered at the beginning of period $i - 2$ to bring recorded inventory back to S. This quantity, RD_{i-3}, will arrive at the beginning of period $i - 1$, but the system will have seen additional demand over the recent period, $i - 2$, given by RD_{i-2}. Thus, once the order of RD_{i-3} arrives, the beginning recorded inventory in period $i - 1$ is given by Eq. 2.

Equation 3 shows how the recorded inventory at the beginning of period i is affected by recorded demand. Since there is a one-period lead time, and orders are placed at the beginning of a period based on the prior period's recorded demand, the recorded inventory at the beginning of period i

will be the incoming order added to the recorded inventory from the end of the prior period.

However, when inventory is removed, but its use is not recorded, the automated replenishment system has no way of knowing where it went, or if it was even removed. Due to this transaction-error inaccuracy, the recorded inventory level will differ from what is physically available or officially on backorder. This is why we define the *actual* net inventory, I. Here, we define actual net inventory at the end of a period, i, given by, I_i. Actual net inventory at the end of period i reflects the depletion of stock by both recorded demand (RD_i) and unrecorded demand (UD_i). This total demand in a period is given by:

$$D_i = RD_i + UD_i \tag{4}$$

We assume that total demand, D_i, occurs according to a Poisson distribution, with rate λ, and we call this day i *actual* demand, denoted $D_i \sim \text{Poisson}(\lambda)$. Since demand is recorded with probability p, we can say that $RD_i \sim \text{Poisson}(p\lambda)$ and $UD_i \sim \text{Poisson}((1-p)\lambda)$.

The above distributions are approximations of the actual distributions of RD_i and UD_i since demand is not recorded in the electronic tracking system when there is no physical inventory on hand. In fact, the automated inventory system is not used to manage backorders in practice, and so the recorded inventory in the system will never reflect backorders. Instead, when a particular item is out of stock, a checksheet is posted near the bin so that nursing staff can record how many items have been needed while the bin was empty. Because we are assuming there is no emergency supply, this checksheet captures backorders (and thus, actual net inventory), and the patient care staff will wait for that particular supply to arrive before using it in patient care (assuming non-emergent care). The material management staff, in practice, will spend extra time placing specific orders for backordered stock - a practice that would help to estimate the cost of backordering. To simplify our approach, however, we allow for recorded demand to potentially register the recorded inventory level as negative as well. In such cases, the model will allow for ordering up to the par level and also satisfy the recorded backorders. Still, the model only applies the costs to backorders based on *actual* net inventory, which will not align with recorded inventory until just after a reconciling count of inventory. We provide a justification for our approximation of the distributions of RD_i and UD_i in the appendix.

Now that we have defined how demand occurs, we briefly mention the handling of incoming orders. We assume that the only inflow of stock to the POU is the incoming replenishment order at the beginning of the period. We assume that there are no errors in the receipt of incoming orders, and so both recorded inventory and actual inventory reflect identically-sized incoming order quantities each period. Therefore, transaction error (failure to record usage) is the only source of disparity between recorded inventory at the end of period i, $\bar{I}'_i$, and actual net inventory at the end of period i, I_i. The following equations help to define the general relationship between actual net inventory at the end of period i, and recorded inventory.
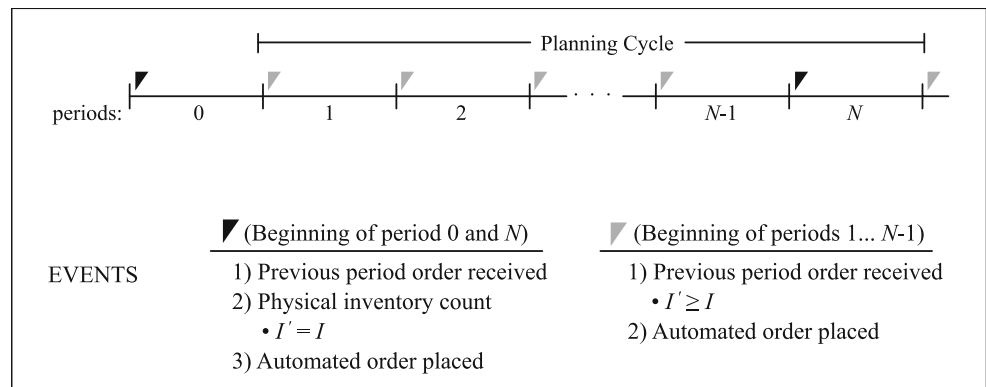
$$I_i = I_{i-1} + (S - \underline{L}'_{i-1}) - D_i = I_{i-1} + RD_{i-2} - D_i \tag{5}$$

$$\bar{I}'_i = \bar{I}'_{i-1} + (S - \underline{L}'_{i-1}) - RD_i = \bar{I}'_{i-1} + RD_{i-2} - RD_i \tag{6}$$

In Eq. 5, we are showing actual net inventory at the end of some period i. Assuming that we know the actual net inventory at the end of period $i-1$, we can compute actual net inventory at the end of period i. Remembering the timing of events (as shown in Figure 3 and as discussed above), we know that the automated system will place an order of size $(S - \underline{L}'_{i-1}) = RD_{i-2}$ at the beginning of period $i-1$, and it will arrive at the beginning of period i. So, the actual net inventory will be increased by this amount exactly, since we have assumed that orders are received without error. And finally, by the end of period i, total demand will have depleted actual net inventory by D_i.

Equation 6 is simply a rewrite of Eq. 1, using the recorded inventory at the end of period $i-1$ as the reference point. Again, the arriving order is RD_{i-2}. However, in the case of recorded inventory, it is reduced by the recorded demand, $RD_i \leq D_i$. So if we assume that both recorded inventory

Fig. 3 Event timeline

| periods: | 0 | 1 | 2 | ... | N-1 | N |

Planning Cycle

EVENTS

▼ (Beginning of period 0 and N)
1) Previous period order received
2) Physical inventory count
 • $I' = I$
3) Automated order placed

▽ (Beginning of periods 1... N-1)
1) Previous period order received
 • $I' \geq I$
2) Automated order placed

and actual inventory are equal at the end of period $i - 1$, then clearly, $I_i \leq \bar{I}'_i$. And in fact, when $p < 1.0$ the hospital found that the relationship was strictly $I_i < \bar{I}'_i$. This difference in values is the source and meaning of inaccurate inventory records, and it leads to greater shortages and an inability to maintain a "stockless" inventory system.

To reconcile the differences between recorded inventory and actual inventory at the POU, a physical inventory count must be performed. We assume that physically counting inventory is an error-free process that does not create additional inaccuracy. We define N to be the number of days between physical inventory counts.

Note in Figure 3 the timing of events over a planning cycle of N periods (days).

Figure 3 shows that counting occurs at the beginning of period 0 after the incoming order is received, and before an order is placed. The next time counting occurs is at the beginning of period N. We will be modeling the daily expected cost over periods 1 through N.

After a count, recorded inventory is reconciled so that $I' = I$ at that instant. Thus, when the order is placed at the beginning of period 0, it is based on accurate inventory records. In other words, an order of size $S - \underline{I}'_0$ is placed when $I' = I$ at that instant. If no demand were to occur during period 0, then the incoming order would bring inventory (both actual and recorded) back up to the order-up-to level, S. Because demand does occur during period 0, *recorded* inventory at the beginning of period 1 will be $S - RD_0$, but *actual* inventory in the beginning of period 1 will be $S - D_0$. We can say precisely what the actual net inventory will be at the end of the first period after a count - simply the beginning period 1 inventory less period 1 total demand:

$$I_1 = S - D_0 - D_1 \tag{7}$$

Subsequent inventory values for actual net inventory at the end of periods $2 \ldots N$ can be derived recursively using Eq. 5.

To apply cost consistently throughout the N periods of the planning cycle, we must define backorders and on-hand inventory at the end of each period. Backorders in a period are based on actual net inventory at the end of the period, and are given by the negative part, I_i^-. Similarly, on-hand inventory is based on actual net inventory at the end of the period, and is simply the positive part, I_i^+. We use the *actual* net inventory for an obvious reason - it represents the actual physical state of the inventory. System inventory costs are applied to the physically available inventory, and that is not what the automated system necessarily reflects.

Having presented the physics of the inventory system, as well as the planning cycle, we now present the cost model formulation.

3.2 Model formulation

In developing the cost minimization model, we now introduce the costs associated with this inventory system. We assume a fixed cost, k, for conducting physical inventory counts. This fixed cost is applied once over the planning cycle of N periods.

As we are assuming full backordering, we define c_b to be the backorder cost per unit per day. Similarly, holding cost per unit, per day, is given by c_h. Both the holding and backordering costs are applied at the end of each period, after demand occurs. Actual net inventory is the source of holding and backordering costs, as recorded inventory may not reflect the physically available inventory or the backlog. We combine Eqs. 7 and 5, applying the relationship given by Eq. 4, to give the following definition for actual net inventory at the end of period i, given some count frequency N.

$$I_i = \begin{cases} S - D_0 - D_1 & \text{for } i = 1 \\ S - \sum_{j=i-1}^{i} D_j - \sum_{j=0}^{i-2} UD_j & \text{for } i = 2 \ldots N \end{cases} \tag{8}$$

You will note that for $i = 2 \ldots N$, actual net inventory is not represented by a recursion as in Eq. 5. The validity of the transformation can be easily shown and we include the derivation in the Appendix.

Then, using the definition for backorders and on-hand inventory, we can give the total cost of the system over N periods:

$$N \cdot C(S, N) = k + \sum_{i=1}^{N} c_h \left[E\left[I_i^+\right] + c_b E\left[I_i^-\right]\right] \tag{9}$$

Equation 9 includes the cost, k, to count once over the cycle N, and the expected holding costs and expected backorder costs that are incurred at the end of each period, i. Because we are assuming that demand occurs according to a Poisson distribution, we can simplify the expectations in the cost equation. Also, we will be minimizing the expected *daily* cost, $C(S, N)$, which will require both sides of Eq. 9 to be divided by N. Simplifying the expectations and averaging over the N periods gives the following daily expected cost objective:

$$C(S, N) = \frac{k + (c_h + c_b) \sum_{i=1}^{N} G^1(S; \mu_i)}{N}$$
$$+ c_h \left(S - \mu - \frac{(N-1)(1-p)\lambda}{2}\right) \tag{10}$$

Recall that total demand $D_i \sim \text{Poisson}(\lambda)$. As demonstrated in Eq. 36 in the Appendix, and based on the relationship defined in Eq. 8, we define Poisson demand rate for

the first period $\mu = 2\lambda$, and after accounting for system-replenished inventory for any period, i, $\mu_i = \mu + (i-1)(1-p)\lambda$.

Similarly, referring to Eq. 38 in the Appendix, we derive the first-order loss function as the expected value of the positive part of demand less the par level, S, $E\left[(\mu_i - S)^+\right] = G^1(S; \mu_i)$, which is the resultant first order loss function for a Poisson distribution with parameter μ_i and par level S. As previously stated, the derivation of Eq. 10 and the development of the new terms can be found in the Appendix.

We now want to find the optimal order-up-to level, S, for a given N. To do this, we take the first difference, over S, of Eq. 10. That first difference is denoted $\Delta_S C(S, N)$ and is given below:

$$\Delta_S C(S, N) = -\frac{(c_h + c_b) \sum_{i=1}^{N} G^0(S; \mu_i)}{N} + c_h \quad (11)$$

Here, $G^0(S; \mu_i)$ is the cumulative cdf for the Poisson distribution, and results from taking the first difference of the Poisson first-order loss function. However, we would achieve the same result for any discrete demand distribution, as we have applied the properties of first order loss functions for discrete ransom variables. Now, we want to find the smallest S that makes $\Delta_S C(S, N) \geq 0$. Therefore, we define S_N^* as the smallest integer such that the following inequality holds:

$$\frac{\sum_{i=1}^{N} G^0(S; \mu_i)}{N} \leq \frac{c_h}{c_h + c_b} \quad (12)$$

Note that the left-hand side (LHS) of the above equation is a non-increasing function of S. So for larger values of S than S_N^*, the inequality will continue to hold. Since the inequality hold is equivalent to Eq. 11 remaining nonnegative, this shows that S_N^* is the global minimum, and not just a local minimum for a particular N.

To find the optimal N^*, we need to explore the optimal solutions, $C(S_N^*, N)$, for a range of values for N. It may be that when $C(S_{N+1}^*, N+1) \geq C(S_N^*, N)$ we have found the optimal solution, with order-up-to level, S_N^*, and counting frequency, $N^* = N$. However, we still need to explore the structure of the model computationally to see whether this is the case.

4 Service level model

In our service level model, we will assume a similar operational structure as described previously. However, as backorder costs are notoriously difficult to estimate, our service level model will forgo backorder costs and instead use a fill rate constraint while minimizing holding and counting costs. We are still assuming an inventory system

for a single item at a single POU location with probability, p, that demand is recorded. N represents the number of days between inventory counts, and S is the order-up-to level. All the parameters are identical, but we do not assume a backorder cost.

In this case, we must define the actual net inventory at the *beginning* of each period, i, denoted NI_i. Since the actual net inventory at the end of some period i is given by Eq. 8, we simply remove the term for period i demand and we get the actual net inventory at the beginning of period i, after the incoming order is received:

$$NI_i = \begin{cases} S - D_0 & \text{for } i = 1 \\ S - D_{i-1} - \sum_{j=0}^{i-2} UD_j & \text{for } i = 2 \ldots N \end{cases} \quad (13)$$

To derive the fill rate for a particular period, we need to compute the number of units short each period, denoted US_i. If we begin a period with outstanding backorders ($NI_i < 0$), then clearly the number of units short will be the demand in that period, D_i. If there is on-hand inventory available, than the number of units short will be the positive part of the difference between demand and available inventory, $(D_i - NI_i)^+$. This idea is formalized in the following equation:

$$US_i = \begin{cases} D_i & \text{if } NI_i < 0 \\ (D_i - NI_i)^+ & \text{if } NI_i \geq 0 \end{cases}$$

$$= \begin{cases} D_i & \text{if } NI_i < 0 \\ \left(D_i - S + D_{i-1} + \sum_{j=0}^{i-2} UD_j\right)^+ & \text{if } NI_i \geq 0 \end{cases} \quad (14)$$

Note the special case when $i = 1$. Following the convention that the summation is zero when the upper limit falls below the lower limit, the relationships still hold true in the model.

From Eq. 14 we can compute the expected number of units short in period i, which turns out to be:

$$\begin{aligned} E[US_i] &= \sum_{k=-\infty}^{S} E[US_i | NI_i = k] Pr(NI_i = k) \\ &= \sum_{k=-\infty}^{-1} \lambda g(S - k; \lambda_i + \lambda) \\ &\quad + \sum_{k=0}^{S} E\left[(D_i - k)^+\right] g(S - k; \lambda_i + \lambda) \\ &= \lambda G^0(S; \lambda(i(1-p) + p)) \\ &\quad + \sum_{k=0}^{S} G^1(k; \lambda) g(S - k; \lambda(i(1-p) + p)) \end{aligned} \quad (15)$$

In the above equation, we have $\lambda_i = (i-1)(1-p)\lambda$, similar to what was used previously. Also, $g(x; y)$, $G^1(x; y)$, and $G^0(x; y)$ represent a pmf, first-order loss function, and cdf, respectively. The y parameter in the previous

statement represents the mean of the respective Poisson distributions for each case, and the x parameter represents the independent variable in each case.

Now that we have the expected number of units short in period i, we can compute the fill rate for period i, denoted $FR(S, i)$. It is simply the compliment of the expected number of units short divided by the expected daily demand, as shown below:

$$FR(S, i) = 1 - G^0 (S; \lambda(i(1-p) + p))$$
$$- \frac{\sum_{k=0}^{S} G^1(k;\lambda)g(S-k;\lambda(i(1-p)+p))}{\lambda} \quad (16)$$

With the fill rate defined for each period, i, we can now address the service level problem. So, instead of using an explicit backorder cost, we define a fill rate constraint. With N days in a cycle, we want the fill rate in the last day of the cycle, $FR(S, N)$ to be at least as much as some specified value, FR_{min}, for example. That constraint is given below:

$$FR(S, N) \geq FR_{min} \quad (17)$$

Notice that the holding costs are increasing in S, therefore we will be minimizing holding costs subject to a minimum service level constraint. So, we choose the smallest integer value of S, denoted as S'_N, that satisfies the service level constraint. Then we simply choose the N to minimize the modified cost objective based on Eq. 10, where we substitute S'_N for S and remove the backorder costs, c_b. This "service level" cost is denoted $C_{SL}(S'_N, N)$, and is given below:

$$C_{SL}(S'_N, N) = \frac{k+c_h \sum_{i=1}^{N} G^1(S'_N;\mu_i)}{N}$$
$$+ c_h \left(S'_N - \mu - \frac{(N-1)(1-p)\lambda}{2} \right) \quad (18)$$

To minimize, we would like to choose the smallest N where $C_{SL}(S'_{N+1}, N + 1) \geq C_{SL}(S'_N, N)$, if it is a unimodal function of N. However, the effectiveness of this approach needs to be explored computationally, as will be seen in a later section.

5 Solution procedure

In this section, we describe the problem instances to be analyzed. We then explore the structure of the cost objective in each model, in relation to the length of the counting cycle, N. Based on the structural findings, we provide solution procedures for each model. In the next section, we explore the physics of the minimizing solutions.

5.1 Problem definitions

For each model we will be considering 891 problem instances. The holding cost, c_h, counting cost, k, demand rate, λ, and accuracy level p, vary identically between both models. The difference between the models is the minimum fill rate constraint, FR_{min}, in the service level model, and the backorder cost, c_b, in the cost-only model. The parameters vary as follows:

$$
\begin{aligned}
c_h &= 0.05, 0.30, 0.60 \\
k &= 20, 40, 100 \\
\lambda &= 8, 15, 20 \\
p &= 0.45, 0.50, 0.55, \ldots, 0.90, 0.95 \\
FR_{min} &= 0.90, 0.95, 0.99 \\
c_b &= 3, 6, 12
\end{aligned}
$$

The values for c_h, c_b, and k used in this model reflect potential magnitudes of differences that might be seen between holding, backordering and counting costs - but do not reflect actual costs in practice. The values for p represent real accuracy ranges seen in practice, with some departments showing accuracy levels as low as 45%, and others as high as 95%, however, the numbers do not represent actual data from the process. Similarly, the demand rates shown reflect the rate at which some products are used in patient care, ranging from as low as 5 units per day to upwards of 30 units per day, depending on the supply item. We believe these arbitrary values we chose represent an acceptable range of values for the purposes of our study - which is to provide insight into how the varying situations and solutions can inform better decision-making within similar environments.

Notice that the parameter for the accuracy level, p, ranges from 0.45 up to 0.95 in increments of 0.05, leading to 11 variations in p. Combined with the three levels of parameter settings for the remaining four parameters, gives a problem set with $11 \cdot 3^4$ instances for each model. Before we can seek solutions for these sample problems, we must explore how the cost behaves in N.

5.2 Structural findings

For both models described in this paper, we have defined optimality conditions for finding S^*_N for a given N and S'_N for a given N and a given fill-rate constraint. However, we need to explore how the related cost objectives for both models behave over changing values of N. If the objectives $C(S^*_N, N)$ and $C_{SL}(S'_N, N)$ are unimodal over values of N, then we can claim that our solutions are optimal using the previously described approaches. Therefore, we will

explore each model computationally to test whether they might be unimodal over N.

After we know the structure of each model, we can proceed with a solution procedure that will minimize the cost objective. For the service level problem we will minimize cost, subject to the fill-rate constraint for period N.

5.2.1 Shortage cost model

To begin our look at the structure of the cost objective, we apply the basic algorithm as described previously. That is, we incremented N until $C(S^*_{N+1}, N + 1) > C(S^*_N, N)$, for each of the 891 problem instances. Then, once we had the potential N^*, we continued to increment N beyond N^* until it was clear that the cost would not decrease below the minimum value. In this way, we explored unimodality while seeking optimal solutions.

The results of this effort showed that by incrementing N until $C(S^*_{N+1}, N + 1) > C(S^*_N, N)$ we were able to find the optimal solution in all but 15 problem instances. The average relative solution distance was just 0.03% above the optimal solution for those 15 non-optimal cases. Furthermore, while we found that the cost function was not unimodal in N for all problem cases, it *was* non-decreasing in our search range beyond N^* for all but 26 problem instances. We cannot say, then, that the function is unimodal in N for all problem instances.

We found that the above method is easy to implement and it also performs quite well, as it found the optimal solution in 98.3% of the test cases and was only off by an average of 0.03% the other 1.7% of the time. Therefore, in the next section we will analyze the results given by using the above method to find optimal or near-optimal solutions to the shortage cost model.

5.2.2 Service level model

For the service level model, we will be looking at how the cost, $C_{SL}(S'_N, N)$, behaves over values of N. Initially, suppose that we set N^* to be the smallest integer such that $C_{SL}(S'_{N+1}, N + 1) > C_{SL}(S'_N, N)$. We want to see how close we can get to the optimal with this simple approach, as with the previous model.

The Service Level model was not found to be unimodal over N, and just over half of the problem instances were found to be non-decreasing after finding N^*. Those that were non-decreasing in our search area found the optimal solution using the simple incrementing approach. For the remaining problem instances, we incremented N beyond N^* enough to see the function turn upwards in a seemingly non-decreasing path. Over this range we took the minimum cost solution to be the optimal.

With this approach, we found that about 74% of the problem instances were able to find optimal solution at N^*. The remaining 26% of the problems had an N^* that was, on average, about 2.4% away from the optimal solution.

As such, the incremental approach performed modestly, but additional work should be done to improve the minimization technique. We leave this additional work for later research, and proceed with our analysis of the results. We hope to gain both structural and managerial insights through our analysis that are useful to both practitioners and academicians.

6 Numerical results

In this section we explore how the parameters affect the minimizing solutions to both the service level problem and the shortage cost problem. As discussed previously, the parameters were largely the same for the separate models, but the shortage costs and fill rate constraints were not present in both models.

6.1 Shortage cost model analysis

As the first model is based on minimizing shortage costs as well as holding and counting costs, we will first look at how the parameters influence cost. We are also interested in the fill rate associated with the different shortage costs, and so we will discuss that later on.

6.1.1 Accuracy level, p

The accuracy level had the same influence on cost as seen in the Regional Hospital's inventory system. When p was greater, the optimal cost was lower. In fact, as p increases linearly from 45% towards 95% accuracy, the cost decreases at a somewhat parabolic descent. If we look at the sensitivity to p, as given by

$$\frac{\Delta C/C}{\Delta p/p}$$

we find that as p increases, the sensitivity to changes in p also increases. Specifically we found that for small p, a 1% increase in the accuracy level equates to only a 0.31% reduction in cost, but when p is large, a 1% increase can mean up to a 3.6% decrease in cost.

Managerially, one would see only minor cost benefit when moving from an accuracy level of, say 45% up to 60%. Once accuracy has been improved sufficiently, the cost is more easily influenced by further improvements in accuracy.

The accuracy level also strongly influences the counting cycle length, N^*. The cycle length seems to increase

exponentially as accuracy increases beyond 80% or 85%, as it is highly sensitive to the value of p. Using the same type of measure as above, we find that for large p, a 1% change in accuracy can lead to an 8% change in N, whereas 1% changes to lesser values of p may only imply 0.45% increases in N.

It is obvious that the sensitivity of the counting frequency to p is the source of the sensitivity in cost. This is an intuitive result, since perfect accuracy would, in theory, require no counting. So, as p approaches 100%, N^* should approach ∞, which would rapidly reduce the cost of counting stock.

The opposite is true for the base stock inventory level. As p increases, we are required to hold less stock. However, since there is a lead time to replenish stock, we are still required to hold stock to meet demand between daily order arrivals, even when accuracy is very high. The data suggest that p influences the base stock level very similarly to the way cost is influenced.

6.1.2 Shortage cost, c_b

The backorder cost, c_b, was found to have less influence on the minimizing solutions than intuition suggests. That is, as the shortage cost increases almost 400%, there is only a slight increase (16%) in the average cost of the minimizing solution. This finding demonstrates that the decision variable solutions are successfully reducing shortages, which was the goal of the model.

The shortage cost had similar but opposing effects on the cycle length, N, and the base stock level, S. The base stock level increases slightly with the shortage cost parameter's increase, while the cycle length decreases slightly. With a higher cost for backorders, it is wise to increase inventory and count more frequently to guard against shortages due to inaccurate records and low on-hand inventory quantities.

While computing cost for the shortage cost model, we also computed the resultant fill rate for each minimizing solution. We found the shortage cost to have a significant influence on the resulting service level, in period N, of the cost-minimizing solution. As the shortage cost increases, so does the fill rate. Unfortunately, there doesn't appear to be a specific shortage cost value in the cost-minimizing model that corresponds to a specific fill rate constraint in the service level problem. Instead, it seems as though a combination of counting frequency, holding cost, and shortage cost values would correspond to specific fill rate constraints.

It is also interesting to note, that the shortage cost value has no influence on how other parameters affect the resulting cost. We find this same type of relationship later when discussing the service level model.

6.1.3 Holding cost, c_h

The analysis shows that holding cost does not greatly influence the results of the minimization. As holding cost is increased by a factor of 12, the cost also increases, by only a factor of 2.5. If we look at cost's sensitivity to c_h, we find that for every 1% increase in holding cost there is only a 0.35% increase in total cost.

An interesting aspect about the holding cost, is that as c_h is increased, both the cycle length, N, and the base stock level, S, are decreased. This means that for higher holding costs, total cost is minimized by counting more frequently and holding less in stock. Basically, this result shows the tradeoff between holding costs and counting costs. If we count more frequently, then we pay the fixed counting cost more often. Therefore, we would reduce holding costs by holding less inventory, which would be safe since counting more frequently ensures that the records maintain enough accuracy to keep shortage costs low.

If we look at the combined influence of holding cost and accuracy level, we find that lower values of holding cost, c_h, lead to a lower cost sensitivity to the accuracy level. That is, for higher c_h, improving the accuracy has a greater affect on cost. This is probably due to the relationship between accuracy and the optimal base stock level, S.

6.1.4 Fixed counting and reconciliation cost, k

As just mentioned for holding cost, an increase in counting costs affects both cycle length and base stock level in a similar fashion. For k, however, the influence is in the opposite direction. So, when k is increased, both S and N are *increased*. Again, this shows the tradeoff between holding cost and counting cost. If it costs more to count, then one would count less and hold more inventory as a result of the decreased accuracy of the ordering system stemming from less frequent reconciliation.

The fixed counting cost also affects total cost in a direct fashion. As k increases, so does the minimized total cost. Although, the overall cost is about as sensitive to k as it is to c_h. This means that a 1% change in k equates to about a 0.3% change in total cost.

We can also see a second-order influence between counting cost and accuracy level. Similar to the relationship between c_h and p, when k is higher, then changes in p have a stronger impact than when k is lower.

6.1.5 Daily demand rate, λ

The Poisson rate for random daily demand was an interesting parameter to explore. It actually had opposite effects on the minimizing solutions of S and N. That is,

as the demand rate increased, S was increased to meet increased demand, while the cycle length decreased to ensure reconciliation happened more frequently.

We also find that as demand rate increases, overall cost is also increased. However, it is interesting that as demand increases, the period N fill rate (for the minimized solution) is actually increased. So as demand increases, the minimizing solution is more likely to alleviate shortages. This result is almost counter-intuitive.

6.1.6 General insights

These results indicate that holding costs, shortage costs, and counting costs can be successfully balanced using the model structure employed here. Managerially, it is wise to have a dynamic cycle length, N, and base stock level, S, based on the accuracy level, p.

Some of the minimizing base stock levels were found to be quite high, and in some cases over 250 units. This amount of stock seems excessive where there is a one period lead time, and where we assume demand has a highest average demand rate of 20 units per day. As a result, a lot of space may be necessary to store certain items in stock. In the hospital setting, where space is valuable and scarce, this may not be feasible. Another type of model, or an adjustment of this model, may be useful to explore for those items which require a lot of space to store the minimizing quantity.

Furthermore, the backorder costs and holding costs chosen in this analysis were arbitrary. A more accurate estimate of actual costs would provide additional insight into this inventory model, and may prove to be more useful to practitioners. We now move on to discuss the service level model.

6.2 Service level model analysis

In analyzing the results of minimizing cost subject to a fill rate constraint, recall that our solution approach found the optimal solutions approximately 74% of the time. As a result, about 26% of our results are near-optimal solutions that were about 2% higher than the optimal solution on average.

We now proceed with our analysis, we use the same pattern as was used in the previous model. We consider the parameters of holding cost, c_h, counting cost, k, demand rate, λ, and the fill rate constraint, FR_{min}, as applied to period N.

6.2.1 Fill rate analysis

It was validated, in every case, that our solution approach did find solutions that met fill rate constraints in every 891 problem instance considered. Remember that the fill rate constraint was applied to the last period of the cycle, period

N, but we are ensured that the resultant fill rates in previous periods (i.e. periods $1, \ldots, N-1$) were higher than the constraint.

In analyzing just the fill rate, only the constraint, FR_{min}, had noticeable influence on the resultant fill rate of the minimizing solution. In fact, there is almost a one-to-one correspondence to between the constraint and the resultant fill rate. This means that the fill rate was a binding constraint in the model. While this seems intuitively obvious, never was there an example where meeting a higher fill rate was less costly than meeting a lower fill rate.

In comparison to the shortage cost model, we find that the fill rate constraint has the same type of effect on cost and minimizing values of S and N as the shortage cost. The difference is that the constrained fill rate generates minimizing solutions with a clear minimum value for the resulting period N fill rate. This is not the case for the resulting period N fill rate from the minimized shortage cost model. In fact, comparing the structure of the fill rate results between models shows significantly different distributions around the average values.

We also found that inventory inaccuracy, counting cost and holding cost were all not shown to influence the resultant fill rate in period N. However, each of the parameters did affect the minimized cost value and the value of the decision variables, S and N, in the minimizing solution. We explore these influences next.

6.2.2 Other parameters: p, c_h, k, λ

Compared to the shortage cost model, the other parameters in the service level model behave nearly identically. For instance, not only does cost increase as accuracy increases, but as accuracy increases, so does the cost sensitivity to the accuracy parameter, as was seen in the previous model. Similarly, the holding cost and counting cost have opposite effects on S and N, as in the previous model. Even the second order effects, which were discussed previously, are found to generate basically the same results in the minimizing solutions.

Similar to the previous model, the resulting base stock levels were quite high, in relation to how a hospital stores and uses inventory. In fact, as in the previous model, there were several cases of the minimizing solution to require a base stock level higher than 200 units. This may not be feasible in a real industry setting, and so there may be some require modifications to the model to address this issue.

7 Conclusion

We develop two healthcare inventory models that allow for full backordering and are based on the use of an

automated inventory replenishment system that is prone to inaccuracy. While the premise of the model is similar to other inaccuracy-related literature, we are the first to introduce both a cost-minimization model and a fill-rate constrained model where demand is random and discrete.

The solution approach we develop was shown to find the optimal base-stock policy for a given counting cycle length. In most test cases, our solution approach also finds the optimal counting cycle length, thus minimizing cost under either shortage costs or fill rate constraints. More specifically, our approach found the optimal solution in 98% of the test cases for the shortage cost model and about 74% of the test cases for the service level model. On average, the near-optimal solutions were off by about 2%.

The models developed here are unique to healthcare management literature, as the unique challenges of the healthcare industry supply chain are underrepresented in the literature. The hospital supply chain considered here operates according to a stockless inventory system, similar to the type defined in [24].

To analyze the effectiveness of the inventory system and of our solution approach, we considered a total of 1782 problem instances, evenly split between the service level model and the shortage cost model. We found that where accuracy, p, is low, there is a crippling affect on the inventory system performance, that is - service level is decreased and cost is increased. We also saw that initial efforts in addressing accuracy (i.e. improving from $p = 0.45$ to $p = 0.70$) will not see as significant results as when improvements are made at higher values (i.e. from $p = 0.80$ to $p = 0.95$). This should be motivation for adopters of a stockless inventory system to strive for only high accuracy values, as the cost savings have been shown to be significant at higher levels of improvement. This should also inform decision-makers on the critical necessity of coordinating with patient care staff to maintain high accuracy in recording inventory.

An application of such findings could likely provide a basis for putting monetary value on inventory accuracy in similar settings in industry. There is currently a thriving industry of inventory system providers that purport to reduce inaccuracy in the inventory system by recording inventory in innovative ways. Having a way to validate these cost savings may strengthen the position of the hospital in making system purchase decisions. Similarly, understanding the model presented here may lead systems providers to assess the effectiveness and value of their own offerings.

The results of the model also show that the base stock inventory system, which is generally used in healthcare, may require a significantly higher base stock level than space allows. In the healthcare industry this would not be feasible, as space is often costly or limited. Previous articles in the literature have explored space limitations in the hospital environment, but do not consider inaccuracy as an obstacle [17, 33]. In a later article currently in progress, we hope to present a model with a modified structure to support lower inventory levels while still keeping costs low.

7.1 Future work

It will also be interesting to study the pervasiveness of inaccuracy in the healthcare industry, as it is already planned as future research. The results of such a study could show how inaccuracy effects both profitability (as in the shortage cost model) and patient care (as in the service level model). As medical supplies are in high demand and are of high value, it is likely that the healthcare industry can see enormous savings in addressing the accuracy issue, as shown in our results and throughout our interaction with industry.

Other future work that applies to what has been done here, may include a modified heuristic approach to increase the possibility of finding optimal solutions. Additionally, we may consider the resultant "hybrid" supply chain that arises when the stockless system cannot be fully implemented, leaving the supply chain network stuck in a kind of operational limbo.

In the research presented here, we employed Maple 12 mathematical programming language to perform our analysis and computation. It may be that other methods of implementing our model my prove to be more efficient.

Appendix

Justification of the approximate distributions used for UD_i and RD_i

The Alabama hospital is using an automated replenishment system that does not record backorders, while the actual net inventory can incur backorders. Because of this, recorded demand, RD_i is dependent on the amount of actual inventory, I, on hand in the POU. Once actual inventory is depleted, demand can no longer be recorded in the system. Similarly, the recorded demand cannot exceed total demand, D_i. In other words, $0 \leq RD_i \leq \min(I, D_i)$. The Figures below demonstrate this idea when $D_i > I$ and $D_i < I$.

When a demand event occurs, and stock is available in the POU inventory location, then the probability that demand is recorded is p. However, if there is no stock on hand when a demand occurs at the POU, the demand is not recorded in the automated ordering system. In practice, the demand is backordered until the next incoming replenishment arrives, leading to a delay in patient care. In the end, the automated system sees neither the backorder, nor its fulfillment. As
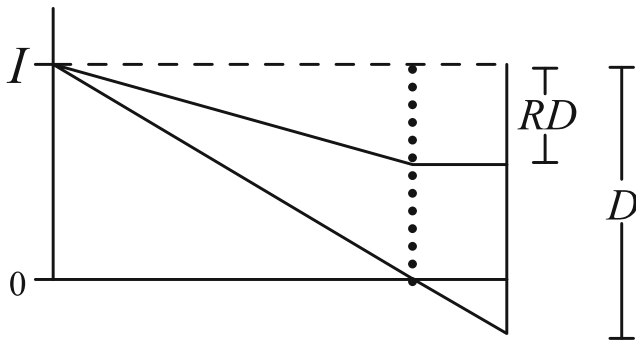
Fig. 4 When RD vs. I

a result, the recorded demand is dependent on the level of on-hand stock and the amount of total demand seen. The figures below demonstrate this concept, where the recorded demand, RD is constrained by either the total demand, D, or the available on-hand inventory, I.

Basically, the recorded demand can be no larger than the smaller of on-hand inventory (Figure 4) versus total demand (Figure 5). This can be written as the following relation:

$$0 \leq RD \leq \text{Min}(D, I) \tag{19}$$

From a demand standpoint, we want to derive the distribution for recorded demand in a particular shift i and day j. Due to the complexity of the actual distribution of recorded demand, we use an approximation of this distribution. We condition on total shift i demand, D_{ij}, which follows a Poisson distribution. Doing this will allow us to derive the approximate distribution of recorded demand, RD_{ij}. The basic distribution equation follows, though we ignore the ij index without loss of generality:

$$Pr[RD = j] = \sum_{k=0}^{\infty} Pr[RD = j | D = k] * Pr[D = k] \tag{20}$$

Since we need to also know the value of available inventory, I, we will be using its value as a reference point as we expand Eq. 20, as follows. Note that we also substitute the Poisson probability as total demand follows a Poisson
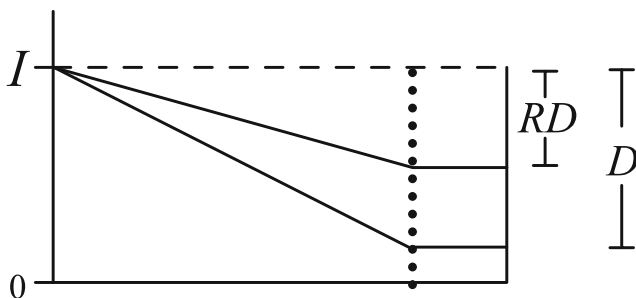

Fig. 5 When RD vs. D

distribution with parameter λ.

$$Pr[RD = j] = \sum_{k=0}^{I} Pr[RD = j | D = k] \frac{e^{-\lambda}\lambda^k}{k!}$$
$$+ \sum_{k=I+1}^{\infty} Pr[RD = j | D = k] \frac{e^{-\lambda}\lambda^k}{k!} \tag{21}$$

Since each demand event has probability, p, of being recorded, we can say that the conditional distribution of recorded demand, RD is simply binomial with probability p. If $D \leq I$ then it will be $Pr[RD = j | D] = \binom{D}{j} p^j (1 - p)^{D-j}$ and if $D > I$, then it will be $Pr[RD = j | I] = \binom{I}{j} p^j (1 - p)^{I-j}$. This relationship affects Eq. 21 as follows:

$$Pr[RD = j] = \sum_{k=j}^{I} \binom{k}{j} p^j (1 - p)^{k-j} \frac{e^{-\lambda}\lambda^k}{k!}$$
$$+ \binom{I}{j} p^j (1 - p)^{I-j} \sum_{k=I+1}^{\infty} \frac{e^{-\lambda}\lambda^k}{k!} \tag{22}$$

Note that the first term in Eq. 22 is in effect when $D = k \leq I$, and the second term when $D = k > I$. The first term can be simplified as follows:

$$\sum_{k=j}^{I} \binom{I}{j} p^j (1 - p)^{k-j} \frac{e^{-\lambda}\lambda^k}{k!}$$
$$= \frac{p^j e^{-\lambda}}{j!} \sum_{k=j}^{I} \frac{[(1-p)\lambda]^{k-j}\lambda^j}{(k-j)!}$$
$$= \frac{(p\lambda)^j e^{-\lambda}}{j!} \sum_{k=0}^{I-j} \frac{[(1-p)\lambda]^k}{k!} \tag{23}$$

Equation 23 can now be written as

$$Pr[RD = j] = \frac{(p\lambda)^j e^{-\lambda}}{j!} \sum_{k=0}^{I-j} \frac{[(1-p)\lambda]^k}{k!}$$
$$+ \binom{I}{j} p^j (1 - p)^{I-j} \left(1 - \sum_{k=0}^{I} \frac{e^{-\lambda}\lambda^k}{k!} \right) \tag{24}$$

Because we cannot condition on I since we do not know its distribution, we consider how Eq. 24 behaves as $I \to \infty$. This will provide us with an asymptotic approximation. The second term will clearly become zero as I approaches infinity. The summation element of the first term has the limit: $e^{\lambda - p\lambda}$ as $I \to \infty$. So, we finally get the following approximate distribution for recorded demand:

$$Pr[RD = j] = \frac{(p\lambda)^j e^{-p\lambda}}{j!} \tag{25}$$

Therefore, we approximate the recorded demand distribution using the Poisson distribution with parameter $p\lambda$.

Remembering the indexing used, we define the recorded demand in shift i of day j to be $RD_{ij} \sim$ Poisson $(p\lambda_i)$. To be consistent with $D_{ij} \sim$ Poisson(λ_i), we see that for unrecorded demand, $UD_{ij} \sim$ Poisson$((1-p)\lambda)$.

In our simulation and our approximate cost versions of the model, we use the above approximations for the distributions of recorded demand and unrecorded demand.

Justification of the actual net inventory equation

$$I_i = \begin{cases} S - D_0 - D_1 & \text{for } i = 1 \\[2ex] S - \sum_{j=i-1}^{i} D_j - \sum_{j=0}^{i-2} UD_j & \text{for } i = 2 \ldots N \end{cases}$$

The above equation is a reprint of Eq. 8 from the text and is derived using the recursive relationship given in Eq. 5 from the text. We want to show that the two forms are identical. First, we look at the case for I_2 and I_3. Then we can prove by induction by assuming true for I_i, $i = 1, \ldots, n$, and then by showing true for I_{n+1}.

We already have that the actual net inventory at the end of the first period is given by $I_1 = S - D_0 - D_1$. At the beginning of period 2 we know that an order will arrive based on the recorded inventory use from period 0, RD_0. Briefly explained, this order of size RD_0 is placed at the beginning of period 1 with a one period lead time. Then, during period 2, actual demand reduces the actual net inventory by D_2. This gives the following recursive equation for the actual net inventory at the end of period 2:

$$I_2 = I_1 + RD_0 - D_2 \tag{26}$$

Substituting the value of actual net inventory from the previous period we have:

$$I_2 = S - D_0 - D_1 + RD_0 - D_2 = S - \sum_{i=1}^{2} D_i - UD_0 \tag{27}$$

which we can say by applying $D_i = UD_i + RD_i$.

Now, for I_3 we have, by the same reasoning:

$$\begin{aligned} I_3 &= I_2 + RD_1 - D_3 \\ &= S - D_0 - D_1 + RD_0 - D_2 + RD_1 - D_3 \\ &= S - \sum_{i=2}^{3} D_i - \sum_{i=0}^{1} UD_i \end{aligned} \tag{28}$$

Now we can see the form given in Eq. 8. We now assume it is true for all periods i from $i = 1, \ldots, n$. This means, that at the end of period $n + 1$, we have the following recursive equation:

$$I_{n+1} = I_n + RD_{n-1} - D_{n+1} \tag{29}$$

The above equation shows the incoming order at the beginning of period $n + 1$ based on period $n - 1$ recorded

demand. It also shows total demand in period $n + 1$, D_{n+1}, reducing the actual net inventory. Substituting the value of period n ending actual net inventory, we have:

$$I_{n+1} = S - \sum_{i=n-1}^{n} D_i - \sum_{i=0}^{n-2} UD_i + RD_{n-1} - D_{n+1} \tag{30}$$

$$= S - \sum_{i=n}^{n+1} D_i - \sum_{i=0}^{n-2} UD_i + RD_{n-1} - D_{n-1} \tag{31}$$

$$= S - \sum_{i=n}^{n+1} D_i - \sum_{i=0}^{n-1} UD_i \tag{32}$$

where we get Eq. 31 by pulling a demand element from the second term of the RHS, and include the demand element from period $n + 1$ in the second term. Then, we apply $D_i = UD_i + RD_i$ to get Eq. 32. And we see that the relationship holds for period $n + 1$, and we have shown that the recursive relationship for end-of-period actual net inventory is identical to Eq. 8.

Justification of the expected daily cost, $C(S, N)$

The notation is the same here as in the main text, though we introduce some new terms. The idea in this section of the Appendix, is to begin by deriving cost for cycle lengths of 1 and 2, as given by $C(S, 1)$ and $C(S, 2)$. From there we use induction to derive cost for any cycle length, N, given by $C(S, N)$, and throughout we develop the optimality conditions for the model solution.

Before we begin, we reiterate some of the assumptions of the model. We are assuming that inventory is counted and reconciled at the beginning of the day prior to the start of the cycle. That means that inventory would have reached S when the order arrived at the beginning of the day 1, but during day 0 there was a total demand seen of D_0. This means that net inventory at the *beginning* of day 1 will always be given by:

$$\underline{I}_1 = S - D_0 \tag{33}$$

With the above starting point, we can now discuss the development of the expected daily cost objective we seek to minimize.

Counting inventory daily

The cost to count inventory every day includes the fixed cost to count, k, added to the holding and backorder costs, both applied to the end of period actual net inventory, I_1. Remember that:

$$I_1 = \underline{I}_1 - D_1 = S - D_0 - D_1 \tag{34}$$

Therefore, when we compute the cost, it will be based on actual net inventory at the end of the period. Note that actual

net inventory is a random variable, which is why we seek to optimize the expected daily cost, $C(S, 1)$, as given below:

$$C(S, 1) = k + E\left[c_h I_1^+ + c_b I_1^-\right] \qquad (35)$$

The above equation gives the expected daily cost, since our cycle is of length $N = 1$ days. To generate the change in daily expected cost over changes in S, we have to first expand the positive part of net inventory into a computable form. But first, we define the following term:

$$\mu = E\left[D_{i-1} + D_i\right] = 2\lambda \qquad (36)$$

where D_{i-1} and D_i are independent and identically distributed Poisson random variables as previously defined for all periods i.

We now have the derivation of expected daily cost in a computable form:

$$\begin{aligned} C(S, 1) &= k + E\left[c_h I_1^+ + c_b I_1^-\right] \\ &= k + E\left[c_h\left(I_1 + I_1^-\right) + c_b I_1^-\right] \\ &= k + c_h E\left[S - \mu\right] + (c_h + c_b)G^1(S; \mu) \end{aligned} \qquad (37)$$

where the first-order loss function is defined as:

$$G^1(S; \mu) = E\left[(\mu - S)^+\right] \qquad (38)$$

Where μ is a Poisson demand rate as given above.

We can define the change in daily expected cost over decision variable, S, when counting inventory daily. The following equation applies the definition of a first difference for a first-order loss function to give the first difference of expected daily cost:

$$\Delta_S C(S, 1) = c_h - (c_h + c_b)G^0(S; \mu) \qquad (39)$$

where $G^0(S; \mu)$ is the ccdf of the Poisson distribution with rate μ and minimum value S. This means that when S increases, the value of $G^0(S; \mu)$ decreases, so that it is a non-increasing function. To optimize the cost function, we form the following inequality:

$$G^0(S; \mu) \leq \frac{c_h}{c_h + c_b} \qquad (40)$$

We then find the smallest S that satisfies the above inequality, and we call it the optimal value for S when counting daily, denoted S_1^*.

We will now take the same approach as we look at counting every other day.

Counting inventory every two days

Now we look at the case when $N = 2$. We still have that inventory is counted and reconciled at the beginning of day 0, so that our starting inventory on day 1 is the same as when we counted daily. Similarly, we still have the same value

for day 1 ending inventory, I_1, but now we need actual net inventory at the end of day 2. This is given below:

$$I_2 = I_1 + RD_0 - D_2 = S - D_1 - D_2 - UD_0 \qquad (41)$$

Since the order placed at the beginning of day 1 is based on recorded inventory, that is the size of the order that arrives at the beginning of day 2. Then, throughout day 2, regular demand occurs which reduces net inventory by D_2. When we expand and apply the relationship $D_0 = RD_0 + UD_0$, we end up with what is given for I_2.

Having a way to describe actual net inventory at the end of day 2, we can now compute the expected cost over a two-day cycle. Recall that we have defined $C(S, j)$ to be the expected daily cost over a cycle of length j. Therefore, we can write the following equation:

$$2 \cdot C(S, 2) = k + E\left[\sum_{i=1}^{2}\left(c_h I_i^+ + c_b I_i^-\right)\right] \qquad (42)$$

From the above equation, we want to now derive the expected daily cost over a cycle of length $N = 2$. The following shows the intermediate steps taken to derive $C(S, 2)$.

$$2 \cdot C(S, 2) = k + E\left[\sum_{i=1}^{2} c_h I_i\right] + E\left[\sum_{i=1}^{2}(c_h + c_b)I_i^-\right] \qquad (43)$$

$$\begin{aligned} 2 \cdot C(S, 2) &= k + E\left[c_h I_1\right] + E\left[c_h I_2\right] \\ &\quad + E\left[(c_h + c_b)I_1^-\right] + E\left[(c_h + c_b)I_2^-\right] \\ &= k + c_h(S - \mu) + c_h(S - (\mu + (1-p)\lambda)) \\ &\quad + (c_h + c_b)G^1(S; \mu) \\ &\quad + (c_h + c_b)G^1(S; \mu + (1-p)\lambda) \qquad (44) \\ &= k + c_h(2(S - \mu) - (1-p)\lambda) \\ &\quad + (c_h + c_b)\sum_{i=1}^{2} G^1(S; \mu_i) \qquad (45) \end{aligned}$$

$$\begin{aligned} C(S, 2) &= k/2 + c_h((S - \mu) - 1/2(1-p)\lambda) \\ &\quad + 1/2(c_h + c_b)\sum_{i=1}^{2} G^1(S; \mu_i) \qquad (46) \end{aligned}$$

To get Eq. 43 we simply apply the definition of $I_i = I_i^+ - I_i^-$, as was done when counting daily, and then combine like terms.

Equation 44 expands the summations for the next step, given by Eq. 44, where we substitute $E[I_2] = E[S - (D_1 + D_2 + UD_0)] = S - \mu - (1-p)\lambda$ as defined by Eq. 41. Note that the values relating to period 1 are the same as when counting daily. Similarly, while the first order loss function is the same for period 1, we introduce $G^1(S, \mu + (1-p)\lambda)$ as the first order loss function for period 2.

To explain Eq. 45, we introduce the following definition which will be used hereafter:

$$\mu_i = \mu + (i-1)(1-p)\lambda \tag{47}$$

Equation 47 defines the Poisson rate at which actual net inventory, I_i, is depleted over the counting cycle.

We substitute (47) into (44) to get Eq. 45. After doing so, we can then solve for the expected daily cost for a cycle of length $N = 2$ days. This is given in Eq. 46. From this point, we can employ the same approach to compute the difference, and we have that:

$$\Delta_S C(S, 2) = c_h - 1/2(c_h + c_b) \sum_{i=1}^{2} G^0(S; \mu_i) \tag{48}$$

Similar to the previous case, we define S_2^* to be the smallest integer value for S that satisfies the following inequality:

$$\frac{\sum_{i=1}^{2} G^0(S; \mu_i)}{2} \leq \frac{c_h}{c_h + c_b} \tag{49}$$

Counting inventory every _N_ days

We now conjecture that the following equation represents the expected daily cost of the inventory system with a counting frequency of N. It is given below:

$$C(S, N) = \frac{k}{N} + \frac{(c_h + c_b) \sum_{i=1}^{N} G^1(S; \mu_i)}{N} \\ + c_h(S - \mu - 1/2(N-1)(1-p)\lambda) \tag{50}$$

We can actually prove the above conjecture by induction. Since we have already shown the above for $N = 1$ and $N = 2$, we will assume that it is true for $N = 1, \ldots, n$, and then show it is true for $N = n + 1$.

We first consider the actual net inventory at the end of day $n + 1$, as given by:

$$I_{n+1} = S - \sum_{i=0}^{n} D_i + \sum_{i=0}^{n-1} RD_i - D_{n+1} \tag{51}$$

which can be explained using the same arguments for period 2. We have seen actual demand, D_i for every period through $n + 1$, and we have seen replenishment arrive based on recorded demand, RD_i, from all periods up to $n - 1$. If we apply the relationship in Eq. 4 we can reduce the above equation to:

$$I_{n+1} = S - \sum_{i=0}^{n-1} UD_i - D_n - D_{n+1} \tag{52}$$

We can then compute the total cost of the individual day $n + 1$ by using the following two formulas as in the previous

cases:

$$c_h E\left[I_{n+1}\right] = c_h \left(S - \mu - n(1-p)\lambda\right) \tag{53}$$

$$(c_h + c_b) E\left[I_{n+1}^-\right] = (c_h + c_b) G^1(S; \mu_{n+1}) \tag{54}$$

Now, we will add the costs for day $n + 1$ to the total cost over n days, $nC(S, n)$. This sum we will denote $(n + 1)C(S, n + 1)$, and we will simplify and combine terms. Then, dividing both sides of the resulting equation by $n + 1$, we have the form we need, given below (completing the proof):

$$C(S, n+1) = \frac{k}{n+1} + \frac{(c_h + c_b) \sum_{i=1}^{n+1} G^1(S; \mu_i)}{n+1} \\ + c_h \left(S - \mu - (n)\frac{1}{2}(1-p)\lambda\right) \tag{55}$$

which proves the conjecture from above.

Knowing the expected daily cost for any counting cycle length, N, allows us to derive the optimality conditions for S_N^*.

References

1. Abdi MMU (2014) Just in Time Procurement and Inventory Management for a Hospital. Master's Thesis for McMaster University
2. Ahmadi E, Masel DT, Metcalf AY, Schuller K (2019) Inventory management of surgical supplies and sterile instruments in hospitals: a literature review. Health Syst 8(2):134–151. https://doi.org/10.1080/20476965.2018.1496875
3. Bijvank M, Vis IF (2012) Inventory control for point-of-use locations in hospitals. J Oper Res Soc 63(4):497–510. https://doi.org/10.1057/jors.2011.52
4. Buschiazzo M, Mula J, Campuzano-Bolarin F (2020) Simulation optimization for the inventory management of healthcare supplies. Int J Simul Model 19(2):255–266. https://doi.org/10.2507/IJSIMM19-2-514
5. Buyurgan N, Zhang S, Okyay HK (2019) Methodical analysis of inventory discrepancy under conditions of uncertainty in supply chain management. Int J Logistics Syst Manage 32(2):272–290
6. Chakrabarty N (2020) A Methodology for Supply Inventory Management for Hospital Nursing Units Considering Service level Constraint, Thesis, Ohio University
7. DeHoratius N, Mersereau AJ, Schrage L (2008) Retail inventory management when records are inaccurate. Manuf Service Oper Manage 10(2):257–277
8. DeScioli D (2005) Differentiating the hospital supply chain for enhanced performance. Ph.D. thesis Massachusetts Institute of Technology
9. Duclos L (1993) Hospital inventory management for emergency demand. Int J Purch Mater Manag, pp 30–37
10. Fleisch E, Tellkamp C (2005) Inventory inaccuracy and supply chain performance: a simulation study of a retail supply chain. Int J Prod Econ 95:373–385

11. Gebicki M, Mooney E, Chen SJG, Mazur LM (2014) Evaluation of hospital medication inventory policies. Health Care Manage Sci 17(3):215–229. https://doi.org/10.1007/s10729-013-9251-1

12. Gurumurthy A, Nair VK, Vinodh S (2020) Application of a hybrid selective inventory control technique in a hospital: a precursor for inventory reduction through lean thinking. TQM Journal (December 2018). https://doi.org/10.1108/TQM-06-2020-0123

13. Hawa T Reimagining the Operating Room Inventory as a Center of Savings by Reimagining the Operating Room Inventory as a Center of (2020)

14. Iglehart DL, Morey RC (1972) Inventory systems with imperfect asset information. Manag Sci 18(8):388–394

15. Jan PJH (2006) Tailored hospital supply chain for greater return on investment. Ph.D. thesis Massachusets Institute of Technology

16. Kök AG, Shang KH (2007) Inspection and replenishment policies for systems with inventory record inaccuracy. Manuf Serv Oper Manage 9(2):185–205

17. Little J, Coughlan B (2008) Optimal inventory policy within hospital space constraints. Health Care Manage Sci 11(2):177–183. https://doi.org/10.1007/s10729-008-9066-7

18. McKone-Sweet KE, Hamilton P, Willis SB (2005) The ailing healthcare supply chain: a prescription for change. The Journal of Supply Chain Management (Winter), pp 4–17

19. Moons K, Waeyenbergh G, Pintelon L (2018) Measuring the logistics performance of internal hospital supply chains - a literature study. Omega (United Kingdom) 0:1–13. https://doi.org/10.1016/j.omega.2018.01.007

20. Morey RC (1985) Estimating service level impacts from changes in cycle count, buffer stock, or corrective action. J Oper Manag 5(4):411–418

21. Opolon DC (2010) Improving product availability in hospitals : the role of inventory inaccuracies. Ph.D. thesis, Massachusetts Institute of Technology, Engineering Systems Division, pp 1–262

22. Raman A, DeHoratius N, Ton Z (2001) Execution: the missing link in retail operations. Calif Manag Rev 43(3):136–152

23. Rekik Y, Sahin E, Dallery Y (2008) Analysis of the impact of the rfid technology on reducing product misplacement errors at retail stores. Int J Prod Econ 112:264–278

24. Rivard-Royer H, Landry S, Beaulieu M (2002) Hybrid stockless: A case study: Lessons for health-care supply chain integration. Int J Oper Prod Manage 22(4):412–424

25. Roper KO, Sedehi A, Ashuri B (2015) A cost-benefit case for RFID implementation in hospitals: Adapting to industry reform. Facilities 33(5-6):367–388. https://doi.org/10.1108/F-05-2013-0041

26. Rosales CR, Magazine M, Rao U (2014) Point-of-use hybrid inventory policy for hospitals. Decis Sci 45(5):913–937. https://doi.org/10.1111/deci.12097

27. Rosales CR, Magazine MJ, Rao US (2020) Dual sourcing and joint replenishment of hospital supplies. IEEE Trans Eng Manag 67(3):918–931. https://doi.org/10.1109/TEM.2019.2895242

28. Rossetti MD (2008) Inventory management issues in health care supply chains. Tech. rep., University of Arkansas

29. Saha E, Ray PK (2019) Modelling and analysis of inventory management systems in healthcare: a review and reflections. Comput Indust Eng 137(November 2018):106051. https://doi.org/10.1016/j.cie.2019.106051

30. Sandoh H, Shimamoto H (2001) A theoretical study on optimal inventory-taking frequency for retailing. J Retail Consum Serv 8:47–52

31. Sheppard GM, Brown KA (1993) Predicting inventory record-keeping errors with discriminant analysis: a field experiment. Int J Prod Econ 32:39–51

32. Upadhyay S, Smith DG (2020) Cash, cash conversion cycle, inventory and covid-19. J Health Care Finance 46(4):37–43

33. Vincent V, Ranton M (1984) Hospital pharmacy inventory management: economic order quantity model with space limitations. Hospital Materials Management Quarterly 5(3):82–86

34. Volland J, Fügener A, Schoenfelder J, Brunner JO (2017) Material logistics in hospitals: a literature review. Omega (United Kingdom) 69:82–101. https://doi.org/10.1016/j.omega.2016.08.004

35. Young ST (1989) Hospital material management: Systems and performance. Journal of Purchasing and Materials Management (Fall), pp 31–35