Journal Pre-proof

Adapting operations to new information technology: a failed “internet of things” application

Michael E. Ketzenberg, Richard D. Metters

PII: S0305-0483(19)30318-4
DOI: https://doi.org/10.1016/j.omega.2019.102152
Reference: OME 102152

To appear in: *Omega*

Received date: 11 March 2019
Accepted date: 5 November 2019

Please cite this article as: Michael E. Ketzenberg, Richard D. Metters, Adapting operations to new information technology: a failed “internet of things” application, *Omega* (2019), doi: https://doi.org/10.1016/j.omega.2019.102152

Highlights

- Optimize inventory routing with stochastic demand, uncertain inventory levels, and emergency pickups.
- Address the value of information through application of “Internet-of-Things” technology.
- IOT application reduces costs by as much as 20.8% and depends on operating conditions.

Adapting Operations to New Information Technology: A Failed "Internet of Things" Application

Michael E. Ketzenberg† and Richard D. Metters

Department of Information and Operations Management
Mays Business School
Texas A & M University
College Station, Texas 77843-4217

Corresponding author, email: mketzenberg@tamu.edu, phone: 979.571.5284

ADAPTING OPERATIONS TO NEW INFORMATION TECHNOLOGY: A FAILED "INTERNET OF THINGS" APPLICATION

(Authors' names blinded for peer review)

Our setting is the secure data destruction industry in which company trucks travel to distant customer sites to pick up outdated documents containing sensitive information (e.g. social security numbers). Outdated documents are stored in locked containers to be shredded, rather than simply recycled, by the secure data destruction service provider. Currently, no client site inventory level information is available. Due to the lack of inventory information, nearly a quarter of client site visits are wasted as there is no inventory to pick up. Another quarter of client site visits arrive too late. To improve operational efficiency, the firm invented an electronic device to make the containers at customer sites a part of the "Internet of Things" - the containers would be able to remotely transmit inventory level information, thus providing the opportunity for just-in-time inventory pickup. Utilizing this information requires a re-structuring of industry operations from fixed scheduled pickups to dynamic routing. Through heuristic creation and robust testing, we examine the conditions in which the Internet of Things application under consideration would be useful or yield minimal gains. Ultimately, the firm decided not to implement the application as the costs of maintenance and the operational re-structuring outweigh potential gains. Our study has implications for firms that are considering to explore Internet of Things to improve their operations.

Key words: Internet of Things; Heuristic; Stochastic inventory routing

1. Introduction

The "Internet of Things" (IOT) refers to objects being sensed or even controlled remotely using an existing network infrastructure. This means objects that formerly operated autonomously are now being networked. From a consumer standpoint, IOT applications allow individuals to start their home air conditioning system from their workplace, and may someday allow their refrigerators to generate a shopping list for them. From a business standpoint, IOT holds the promise of machines that continually tell us when they are wearing down, replacing expensive and infrequent manual inspection. More relevant to this effort, IOT holds the promise of accurate information about remote objects. For a full classification of IOT, we refer the readers to Atzori et al. (2010). In this study, we examine the efficacy of a specific business IOT application that is invented by a secure data destruction service provider (SDDSP) to improve operational efficiency in the secure data destruction industry.

The secure data destruction industry has arisen to serve the need for secure, reliable document destruction. Firms are responsible for protecting customers and employees from having personal, sensitive data fall into the wrong hands. Often, sensitive data such as social security numbers, medical records, banking or legal records are printed on paper. When these sensitive documents are no longer needed, they cannot merely be recycled, they must be securely destroyed (shredded) prior to recycling, a service typically outsourced to SDDSPs. Hence, SDDSPs pick up outdated sensitive and confidential documents from companies to destroy and protect those documents until they do so.

The emergence of the secure data shredding industry dates back to 1950s. Scandals such as Watergate, Iran-Contra, Enron, and Worldcom have collectively increased the need for secure, reliable document destruction (Wagner 2008). Beyond the ethical obligation, regulation to combat the growing problem of identity theft is another factor driving industry growth (Gutek and Shroff 2005). Among them, the Economic Espionage Act (1996), Health Insurance Portability and Accountability Act (HIPAA, 1996), Gramm-Leach-Bliley Act (GLBA, 1999), and Fair and Accurate Credit Transactions Act (FACTA, 2003) are representative examples. Similarly, there are also state-specific enacted identity theft legislation. The increasing concern to protect confidential information is fueling industry growth. The "document management services" is estimated to be a $6.5 billion market in the U.S. alone (IBISWorld 2019).

From an operations standpoint, "business as usual" in this industry is quite simple in concept. A typical arrangement with a client calls for a specified number of secure containers provided by the SDDSP to be at the client's site. The SDDSP visits the client on a regular schedule (e.g., once per week) and empties the containers. However, "business as usual" has an operational problem and a strategic problem. Operationally, the SDDSP has no way to determine how many documents are accumulated in the containers. The containers are made to be secure to make sure no one, not even client firm employees, could see or physically steal sensitive data. The containers are locked, with only the SDDSP having keys, they are non-see through, and have only a small slit in the top to put in paper. Even the customers who are on site have no idea how full the container is until it nears overflowing. The containers, being aesthetically unappealing, are often hid under counter-tops or in closets, further preventing client firm employees from ascertaining inventory levels. Since the accumulation of sensitive documents is stochastic, there is the classic trade-off that occurs in most fixed-time inventory systems: (1) visiting too frequently means containers are often empty and trips are wasted; and (2) visiting too infrequently means containers fill long before the next scheduled visit, causing a potential security breech, customer displeasure, and costly emergency pickups. Of our 1,008 observations of customer site inventory, on 23% of pick-ups the containers were less than 10% full. On 24% of pick-ups, the containers were over 90% full or overflowing. The lack of container status information generates the difficulty in choosing the right frequency, making the visits operationally inefficient. Strategically, there is a fierce competition in the industry. Since there are virtually no barriers to entry, anyone with a truck and a paper shredder could start a shredding business and easily target the customer base of major service providers. The SDDSP we study estimates that 20% of the industry is already composed of those "mom and pop" organizations.

To improve operational e□ciency and gain a distinct advantage over competitors, the Information Technology group of the SDDSP we study invented a new prototype device that could be attached to the inside of their collection containers so that the "dumb" containers could join the IOT. In e□ect, the technology enables the service provider to track the container inventory status at each customer facility. This device measures the amount of space left in the container with an optical infra-red sensor and transmits the data over the internet to the SDDSP. Note that it would seem the technologies involved here are not new. However, the specific environment militates against existing equipment. Since the device is located in a mobile container, it cannot be plugged in, so the device has to have its own power source. It has to be physically tough enough to withstand the container unloading process without damage. Lastly, it has to be compatible for mounting and installing on the containers by the existing labor force.

Even though the SDDSP was able to obtain real time inventory status information with the new IOT application, the use of such information was not clear to the firm. To utilize the real time container fill level information, the static "once per week" pick-up schedule would need to be scrapped in favor of a dynamic pick-up schedule. However, the Operations group did not know how to create a dynamic schedule. Dynamic routing is not as simple as "just pick up full containers". For example, if a half-full container is three blocks away from the full container that must be picked up, do you pick up the half-full container while you are close by, or wait until later? Further, even though the production and the maintenance cost of the new IOT application was known, the operational benefit was not clear. Consequently, the authors were called in to help the firm assess the value of the new IOT application to decide whether the firm should engage in production and installation of multiple thousands of units.

Common sense generally assumes that "more information is better", and one would assume that the new IOT application should provide significant savings to the SDDSP. Consistent with this intuition, Smartbin, a leading global supplier of IOT ready containers, claims that the recycling market, including secure data shredding, in Europe wastes up to $1 billion a year on deploying ine□cient practices due to a lack of information (Smartbin 2009). Further, the company estimates that service providers incur over $1,000 per container per year of additional costs through the sheer ine□ciency associated with the unnecessary and costly collections of containers that are only partially full (Smartbin 2010). Even so, these and other similar marketplace claims are unsubstantiated. We are aware of no research that quantifies the value of the this type of IOT application and identifies the drivers of value. Our research addresses this gap in the literature.

The usual procedure for assessing the value of the new IOT application should be to compare the cost of "business as usual" to the cost of the business with IOT. Ideally, such a comparison should be based on the full data of all the company's operations. However, there are two factors that militated against this. First, the firm believed it had data but that belief was illusory. For instance, route drivers provided inaccurate, exaggerated data on time requirements to allow themselves to have more leisure time during the day. They also provided greatly exaggerated estimates of the percentage of container fullness, fearing loss of work if they reported that they were "emptying" already empty containers. Consequently, one of the authors accompanied a driver in the Atlanta, USA o□ce for 16 days to gather data to support the parameter estimation in our analysis.

Second, the SDDSP has no structured operational policy for "business as usual". Rather, it has a collection of agreements to pick up containers at customer facilities on specific days. The method to convert these agreements into operational route decisions is ad hoc, produced at each local site by a dispatcher based on gut feel with a goal of maximizing visits rather than minimizing cost, and constantly changed based on customer churn. Comparing any systematic policy developed for the business with IOT to the SDDSP's current policy is not appropriate since such a comparison would considerably and unfairly overestimate the value of the new IOT application.

To assess the value of the new IOT application properly, we first constructed a well performing algorithm that can achieve the best operational outcome without the new IOT application for a static, cyclic environment. Next, we constructed a dynamic vehicle routing algorithm that uses the real time container fill level information from the IOT application and developed a dynamic pick-up schedule. (It should be noted that other possibilities exist: for example, an anonymous referee suggested various hybrid policies that partially use sensors.) For both static and dynamic cases, routes are designed for a single vehicle. We formulate the model as a Markov decision process (MDP). Demand is stochastic. Customers are not satisfied if containers are not serviced in a timely manner. If a container is full and not scheduled to be serviced, documents are at risk. If that customer is not scheduled on a route, the service provider sends a separate truck in an "emergency" fashion to alleviate the problem by emptying the container as soon as possible. The objective is to determine a routing and pickup frequency for customers that minimizes the expected average total cost per period over an infinite horizon. Total cost is composed of two parts: inventory cost (including a penalty cost for emergency pickups, a loss of goodwill cost, and a holding cost) and transportation cost (including a truck setup cost, a transportation cost between facilities on a route, and a fixed collection cost).

Given the complexity of the MDP, we provide some analytical insights to the problem under restricted assumptions. Not surprisingly, the unrestricted problem is too complex for solving to optimality, so we created heuristics for both schedules. The static, cyclic routing heuristic achieves a cost performance that is on average 1.4% higher than optimal. Our experimental results indicate that the value of the IOT application ranges between 0.0%-20.8%, with an average of 5.7%, depending on the parameter settings. We also show that the value increases as the geographic dispersion of customer locations increases and as route capacity decreases. Ultimately, we found that the IOT application is only helpful for some of the firm's customers, while for others the value it generates is quite low. The devices had maintenance cost issues: the containers are mobile and not connected to an electric grid, so the devices had to be powered by AAA batteries, which would need frequent replacement. Further, moving away from static, cyclic customer visitation to a dynamic schedule caused discomfort. Consequently, the SDDSP decided to abandon the project.

While we specifically study the industry of secure data destruction, this type of IOT application adoption, the basic trade-offs, and the inherent change to operations to adapt to the new technology are also present in many other situations where firms must visit remote, unmanned sites in which the status of inventory is uncertain. These include, but are not limited to, propane/oil delivery (Bell et al. 1983, Dror and Ball 1987), vending machines (Ketzenberg et al. 2013), waste removal (Faccio et al. 2011), recycling pickup (Prive et al. 2006), medical waste collection (Birpinar et al. 2009), and oil/gas well monitoring. Therefore, our approach in this study can be more broadly applied to other industries even if it is motivated specifically by the secure data destruction industry.

Despite the prevalence of the problem, there is little research that addresses the unique problem faced by the SDDSP industry. Even if we generalize the context of our study to include different vendor managed inventory (VMI) systems, surprisingly, we find no research that addresses an inventory routing problem (IRP) where (1) customer inventory levels are *unknown* prior to route decision making, (2) customers request emergency pickups when demand exceeds capacity, and (3) demand is stochastic. In fact, it is interesting to note that most contributions to the IRP literature assume that inventory levels are known (e.g. Minkoff (1993), Kleywegt et al. (2002)). This is a clear disconnect with practice since in most situations inventory levels are not known. Hence, we contribute from a modeling perspective by formulating an IRP with these three factors.

The rest of this paper is organized as follows. In §2, we review literature. In §3, we discuss the characteristics of the problem under consideration and introduce the policy without the IOT application and the policy with the IOT application. In §4, we introduce the heuristic. In §5, we assess the value of the IOT application with a numerical study under different parameter settings and routing characteristics. Finally, we conclude the paper with comments in §6.

2. Literature Review

Our research overlaps the two literature streams of IRP and value of information (VOI). In this section, we provide a brief review of the literature in each area and then position our research at their intersection.

2.1. Inventory Routing Problem Literature

The IRP involves the integration and coordination of two components of supply chain management: inventory management and vehicle routing. Essentially, the problem concerns the distribution of one or more products from a single facility to a set of n customers over a planning horizon that may be finite or infinite. The general objective of the IRP is to minimize transportation and inventory related costs over the planning horizon.

The IRP literature has grown considerably since the pioneering papers that appeared in the early 1980s (Bell et al. 1983, Federgruen and Zipkin 1984, Dror et al. 1985). Coelho et al. (2014), and Andersson et al. (2010) provide broad overviews of the IRP literature since that time. The models in the IRP literature can be classified into two groups based on whether demand is deterministic or stochastic. Our paper belongs to the latter group. When demand is stochastic and the planning horizon is infinite, formulating the IRP as a MDP is common and is the method we follow here. In the IRP literature, the general assumption is that the inventory levels are known. Minkoff (1993), Kleywegt et al. (2002, 2004), and Hvattum et al. (2009) are examples of this kind of work. However, no paper in the IRP literature quantifies the value of knowing this information and using it in decision making. Moreover, none of these papers considers emergency pickup/delivery.

We assess the value of a new IOT application in a multiple delivery problem in which a single vehicle can visit one or more customers on a route. In addition, it is also possible to send a vehicle on an emergency pickup (a.k.a. direct delivery in the literature) if demand exceeds capacity. As will be described operationally later, emergency pickups are an important aspect of the secure data destruction industry. Emergency pickup/delivery is also common in other industries, but has not received much attention in research. Jaillet et al. (2002), Chen and Lin (2009), and Hemmelmayr et al. (2010) are some of the papers that address emergency pickup/delivery. Jaillet et al. (2002) formulate an IRP for a finite time period. Even though they assume stochastic demand, daily demand is capacitated. They also assume that the cost for each scheduled delivery at a given customer is fixed and does not change regardless of how many other customers are also visited on the same route. Chen and Lin (2009) study an IRP that involves repeated delivery of multiple products over an extended period. They allow emergency delivery in case of a shortage between two review period occurs. However, the occurrence of emergency delivery does not change the review cycle. Moreover, neither of those papers addresses the direct impact of emergency pickup/delivery on the optimal routing policy.

Hemmelmayr et al. (2010) address the problem of delivering perishable blood products to hospitals. In this case, a vendor prevents blood product shortage at hospitals with different emergency delivery alternatives. One of the alternatives they propose is very similar to the emergency pickup practice we model in our paper. However, they only address a transportation cost minimization problem whereas, in addition to transportation cost, we also consider other costs that include setup, collection, penalty, and holding costs. In addition, we note that none of the papers mentioned here that incorporate emergency pickup/delivery address the value of an IOT application.

2.2. Value of Information Literature

In the supply chain management literature, the VOI has generally been explored with respect to inventory management and not in conjunction with routing decisions. The VOI literature can be classified into two groups based on the information shared: downstream information and upstream information. We position our paper with respect to the contributions addressing downstream information since that is the focus of our study.

The literature addresses several types of downstream information-sharing. Among them demand or inventory data (Cheung and Lee 2002, Ketzenberg et al. 2013), order policy (Ketzenberg and Ferguson 2008, Simchi-Levi and Zhao 2003), advance order information (Zhang and Cheung 2011, Hariharan and Zipkin 1995), projected future orders (Chen and Lee 2009), collaborative forecasting or shared forecast (Aviv 2002, Mishra et al. 2009), product return information (Ruiz-Benitez et al. 2013) and planned promotion information (Iyer and Ye 2000) are all representative examples.

The majority of the literature discussed above study the VOI for managing inventory. Among them, Ketzenberg et al. (2013) is the closest to our own since that paper, like ours, addresses the VOI for a VMI setting in which inventory is monitored remotely. They study the problem in the context of the vending industry. Similar to the sensors in secure data containers, technology allows the supplier to remotely monitor the level of inventory in a vending machine. They develop an analytical model with and without information for the supplier to determine the optimal replenishment cycles for vending machines and quantify the benefit of this technology. They show that the VOI is not straightforward and depends on certain model parameters. However, they ignore routing since the vending machines are located close to each other and therefore transportation is not a significant issue. In contrast, we address operating environments in which customers are geographically dispersed and transportation costs are significant.

The VOI in the context of coordination of inventory management and vehicle routing has not received much attention in research. We are not aware of any research that addresses the VOI from multiple, geographically dispersed inventory locations to dynamically schedule routes when there are significant transportation costs. Cheung and Lee (2002) and Ruiz-Benitez et al. (2013) are, however, closely related papers to our own that explore the VOI with a focus on reducing transportation cost. Cheung and Lee (2002) model a supplier serving multiple retailers. The retailers are located in close proximity so that all the retailers are resupplied in one single delivery. Therefore, they ignore the inter-retailer transit times and do not consider route decisions. The supplier uses retailers' inventory position information to coordinate shipments and re-balance their stocking positions.

Ruiz-Benitez et al. (2013) study a control policy for a manufacturer to manage the collection of returned products from retailers. The decision of interest in that study is the frequency in which returns are picked up from a collection point. Since the value of the returns decay over time, there is a tradeoff between reducing transportation costs with less frequent visits and increasing the value of asset recovery with more frequent visits. The VOI in this context is measured as the reduction in cost that arises from knowing the exact quantity of returns at the collection point, relative to the case of only using the historical distribution of returns. Since a single collection point is modeled, routing is not addressed and emergency pickups are not considered.

To the best of our knowledge, our paper is the first in the literature that evaluates the value of a new IOT application for the coordination of inventory management and vehicle routing.

3. The Model

In this section, we first discuss the business environment to provide support for major modeling assumptions. Then, we introduce our model framework and construct an MDP to determine an optimal routing schedule for both the base policy and the information policy. Finally, we develop an analytical expression for a special case that provides limited insights to the problem and inspires our heuristics developed to analyze the value of the new IOT application.

3.1. The Business Environment

The general setting involves a SDDSP that schedules routes to collect sensitive documents for destruction from N customer facilities using a single vehicle. The service provider then sells the shredded documents to an exogenous buyer. Our SDDSP uses 13 vehicles based in the city of Atlanta to serve the state of Georgia. There is a long tradition in the routing literature of the "cluster-first route-second" approach (e.g. Fisher and Jaikumar (1981), Yang et al. (2000)). The main idea with this approach is that customers are grouped together as a first step, then routing decisions are made on each customer group. This reduces an M-vehicle problem to M, one-vehicle problems. This approach makes particular sense for our SDDSP for three reasons. Primarily, the customers are widely distributed geographically, creating obvious customer groupings. Each black dot on Figure 1 represents a customer location. The four circles correspond to four distinct geographic set of locations, each assigned to a single vehicle, and thereby accounting for four of the 13 vehicles in this facility. Secondly, customer idiosyncracies militate against new/different drivers: our drive-alongs with the SDDSP demonstrated that the containers are frequently hidden from view, often behind unnumbered doors. A driver must know where exactly they are. Facility access is not always obvious, nor is temporary parking for the truck clear so experience saves time. Thirdly, given the sensitive nature of the documents, drivers must pass through security to get to them. Known drivers pass through quickly. Even the presence of the authors on our ride-alongs frequently presented an obstacle, as several firms delayed us at security and made phone calls to verify our status, despite the fact that we were with their trusted driver. In short, while it is possible to switch customers between routes, it would be time consuming and impractical to do so. As a consequence, we adopt the practice of Fisher and Jaikumar (1981) by clustering customers and sequentially solving a series of one-vehicle problems.

Figure 1 **Representative customer locations**

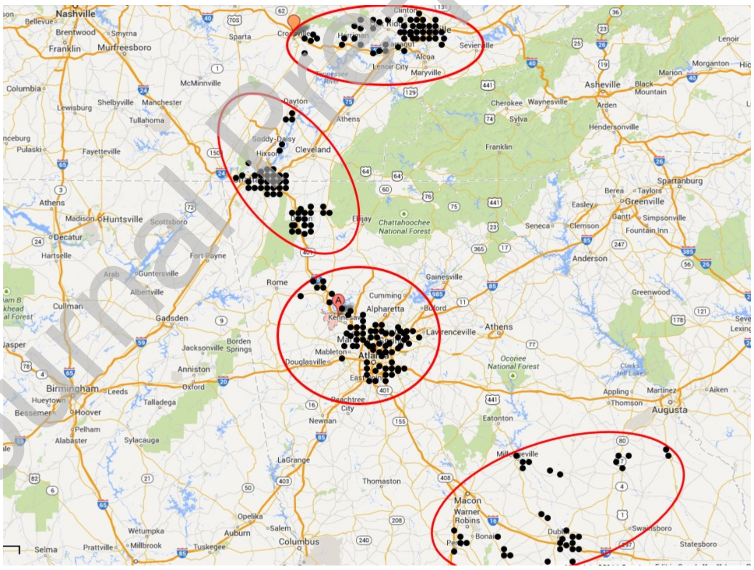

We assume that vehicle capacity is not a consideration. The vehicle fleet was purchased so that vehicles had the capacity to handle the largest customers. For example, bank check processing centers must destroy all checks after the information is electronically captured, filling several dozen 320 gallon containers daily. A dedicated truck is sent to these customers and is solely at that customer site all day. The customers that require a routing solution and non-daily pick-up usually have one or two 42 gallon containers. Drivers must service these customers during normal business hours. Consequently drivers run out of time long before they run out of vehicle paper holding capacity.

Let i denote a customer, $i \in \{0,1,2,\ldots N\}$, where the special case of $i = 0$ denotes the service provider itself. The disposal of documents at a customer represents the demand for the service provider. Sensitive documents accumulate in secured containers provided by the SDDSP that are emptied periodically. The IOT containers include advanced sensors and intelligent processors that allow the service provider to know the container status prior to making a route decision. It was decided by the firm that only daily bursts from the IOT containers would be used, as continuous information would exacerbate battery drain.

In each period t, the decisions of interest for the service provider are the set of customers to visit and the order in which to visit within the selected set. The order of events in each period is (1) select a set of customers to visit and determine a route based on (i) the number of periods since the last visits to all customers for the base policy and (ii) the container status information of all customers for the information policy, (2) visit the scheduled customers and realize demand (empty containers), and (3) visit unscheduled customers that need emergency pickup. To keep the problem tractable, we assume that the service provider is perfectly reliable. All scheduled customers are serviced without fail.

3.2. Cost Structure

We assume that containers are identical across customers and have a capacity of C. Containers have different capacities in practice. A specific container size is chosen by the customer based on their aesthetic preferences, demand rate, and size of the available collection area. Nevertheless, a difference in container size does not change container transportation or emptying cost. Moreover, we allow demand rates to vary across customers. Containers are emptied by an automated mechanism at the truck, so both light and heavy loads take the same amount of time. A fixed collection cost K is incurred during a customer visit that corresponds to the labor involved. We also assume a truck setup cost F that represents the labor cost of employing a driver to prepare a truck to run a route. There is one driver per vehicle and one vehicle per route.

The transportation cost from customer i to customer j, c_{ij}, is known and proportional to traveling time. As is typical in the IRP literature, we assume that $c_{ij} = c_{ji}$. All customers on a route must be serviced in less than H hours for the route to be feasible. This constraint reflects the issue that customer facilities must be accessed within their operating hours. Let $d_t = 1,2,3,\ldots,k$ denote a feasible route decision at time t, where k denotes the number of all feasible routes. From here forward, we suppress the subscript t when the context is clear. Hence, $d \equiv d_t$.

There is an opportunity cost of not collecting the sensitive documents since the service provider delays selling recycled paper from the shredded documents it collects. We represent this opportunity cost as a holding cost. Specifically, we assume a holding cost h per unit due to the uncollected documents, where a unit corresponds to a fixed percentage of container capacity.

If the container fills prior to be serviced, sensitive documents that do not fit in the container may lay around where unauthorized people have access to them. Therefore, customers are dissatisfied and the service provider incurs a loss of goodwill cost p. When a container fills, the customer requests an emergency pickup. If that customer is scheduled on a route, the container is emptied by the scheduled truck. If the customer is not scheduled on a route, the service provider sends a separate truck for an emergency pickup in a "direct delivery" fashion and empty the container. For this emergency pickup service, the service provider incurs an emergency pickup cost P *in addition to* the loss of goodwill cost p. In reality, P is some function of an expediting cost, plus F, K, and c_{0i} and hence can vary from one customer to another. However, this would unnecessarily complicate the model since the fixed costs overwhelm the variation we might observe across customers.

3.3. Business without IOT Application

We call the policy constructed without the IOT application as the base policy. The service provider does not know the amount of paper in the containers before the IOT application. However, it can track the number of periods since the last visit to a customer. With this information, the container status can be estimated given a known probability distribution of demand at each customer facility. Let x_{it} denote the number of periods since the last visit to customer i at time t where x_{it} $\{1, 2, 3, 4, \dots\}$. Hence, accumulated documents in a container become a function of x_{it} for customer i at time t. Then, $\vec{x}_t = [x_{1t}, x_{2t}, \dots, x_{Nt}]$ denotes the state of all customers at time t. After observing the state of all customers, the SDDSP decides which customers to visit. Let a_{it} be a binary indicator variable that defines whether ($a_{it} = 1$) or not ($a_{it} = 0$) customer i is visited on route d at time t.

Let U_{it} denote the demand realization of customer i at time t. The demand for storage space at each customer i is random and i.i.d. over time with mean μ_i, probability mass function (pmf) $\phi_i(\cdot)$, and cumulative distribution function (cdf) $\Phi_i(\cdot)$. Now, let $\phi_i^x(\cdot)$ and $\Phi_i^x(\cdot)$, respectively denote the pmf and cdf for the x-fold convolution of the demand distribution. Then, $U_i^x = \sum_{z=t-x_{it}+1}^{t} U_{iz}$ represents the inventory level at time t. Note that an emergency pickup for customer i at time t arises when $a_{it} = 0$ and $U_i^x > C$. We assume that any demand realization occurs before the driver arrives at the customer facility.

For any state x_{it}, a_{it}, and U_{it}, the possible states in the next period for customer i are given by $x_{i(t+1)}$ where

$$x_{i(t+1)} = \begin{cases} x_{it} + 1 & \text{if } a_{it} = 0 \text{ and } \sum_{z=t-x_{it}+1}^{t} U_{iz} \leq C \\ 1 & \text{otherwise} \end{cases} \tag{1}$$

Equation 1 shows that if a customer is not visited (the first expression), the number of periods since the last visit increases by one period. If the customer is visited on a route or due to an emergency pickup (the second expression), the state for that customer in the next period will be one, implying that it has been one period since the last visit.

Now, let $\omega^B(x_{i(t+1)} \mid x_{it}, a_{it})$ denote the one period transition probability in the base policy for customer i from state x_{it} to state $x_{i(t+1)}$, given both x_{it} and a_{it}. Hereafter, the superscript B denotes the base policy. Additionally, let $\eta^B(\vec{x}_{t+1} \mid \vec{x}_t, d)$ denote the one period joint transition probability from state $\vec{x}_t$ to state $\vec{x}_{t+1}$, given that the decision d has been implemented. Therefore

$$\omega^B(x_{i(t+1)} \mid x_{it}, a_{it}) = \begin{cases} 1 & x_{i(t+1)} = 1 \text{ and } a_{it} = 1 \\ \Phi_i^x(C) & x_{i(t+1)} = x_{it} + 1 \text{ and } a_{it} = 0 \\ 1 - \Phi_i^x(C) & x_{i(t+1)} = 1 \text{ and } a_{it} = 0 \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

and $\eta^B(\vec{x}_{t+1} \mid \vec{x}_t, d) = \prod_{i=1}^{N} \omega^B(x_{i(t+1)} \mid x_{it}, a_{it})$.

Note that when the SDDPS decides to visit customer i at time t (i.e. $a_{it} = 1$), the state for customer i in the next period will be 1 with probability of one. When customer i is not scheduled on a route (i.e. $a_{it} = 0$), there are two possible scenarios. First, if the inventory level does not exceed the container capacity (i.e. $U_i^x \leq C$), an emergency pickup for customer i does not arise and the state in the next period becomes $x_{it} + 1$. This scenario occurs with a probability of $Pr(U_i^x \leq C) = \Phi_i^x(C)$. Second, if the inventory level exceeds the container capacity (i.e. $U_i^x > C$), an emergency pickup for customer i arises in the current period and the state for customer i in the next period will be 1. This scenario occurs with a probability of $Pr(U_i^x > C) = 1 - \Phi_i^x(C)$. Hence, equation 2 captures all possible scenarios.

Since we are not interested in transient behavior or end of horizon effects, we formulate the IRP as a MDP where the objective is to find the optimal routing policy so that the service provider's long-run average expected cost per period is minimized. The linkage between periods is captured through $\eta^B(\vec{x}_{t+1} \mid \vec{x}_t, d)$.

Given the starting state $\vec{x}_t$, the one period loss function $L^B(\vec{x}_t, d)$ that is composed of the expected emergency pickup cost P, the loss of goodwill p, and the holding cost h is

$$L^B(\vec{x}_t, d) = \underbrace{P \sum_{i=1}^{N} \left[(1 - a_{it}) \cdot \sum_{U_i^x > C} \phi_i^x(U_i^x) \right]}_{\text{emergency pickup cost}} + \underbrace{p \sum_{i=1}^{N} \sum_{U_i^x > C} \phi_i^x(U_i^x)}_{\text{loss of goodwill cost}} + \underbrace{h \sum_{i=1}^{N} \left[(1 - a_{it}) \cdot \sum_{U_i^x \leq C} U_i^x \cdot \phi_i^x(U_i^x) \right]}_{\text{holding cost}} \tag{3}$$

Let TC_d be the total traveling cost $\sum c_{ij}$ of route d. Now, let s_d denote the number of customers on route d. Then, the one period total transportation function that is composed of the truck setup cost F, the total traveling cost TC_d, and the fixed collection cost K is given by $T(d)$ where

$$T(d) = \begin{cases} \underbrace{K \sum_{i=1}^{N} a_{it}}_{\text{collection cost}} + F + TC_d & \text{if } s_d \geq 1 \\ 0 & \text{if } s_d = 0 \end{cases}$$

Given the starting state $\vec{x_t}$, the cost to go, if future periods behave optimally, is $f^B(\vec{x_t}, d)$. As is the custom in average cost dynamic programming models (Ketzenberg and Ferguson 2008), we use $\bar{z}$ to denote the equivalent average cost per period when an optimal policy is used. The service provider's objective function in the base policy is given by

$$\pi^B = \min_{1 \leq d \leq k} f^B(\vec{x_t}, d) + \bar{z} \tag{4}$$

where $f^B(\vec{x_t}, d) + \bar{z} = T(d) + L^B(\vec{x_t}, d) + \sum_{x_{t+1}} f^B(\vec{x}_{t+1}, d) \cdot \eta^B(\vec{x}_{t+1} \mid \vec{x_t}, d)$.

3.4. Business with IOT Application

We call the policy constructed with the IOT application as the information policy. In the information policy, the service provider makes route decisions based on the container status information obtained from sensors at each customer facility. Let I_{it} denote the container status at customer i at the beginning of period t where I_{it} $\{0, 1, 2, \ldots, C\}$. Any excess paper beyond the capacity of the container is "backlogged" as the service provider empties a full container in the worst case with an emergency pickup. Then, $\vec{I_t} = [I_{1t}, I_{2t}, \ldots, I_{Nt}]$ denotes the container status of all customers at time t.

Note that an emergency pickup in the information policy for customer i at time t arises when $a_{it} = 0$ and $I_{it} + U_{it} > C$. For any state I_{it}, a_{it}, and U_{it}, the possible states in the next period for customer i are given by $I_{i(t+1)}$ where

$$I_{i(t+1)} = \begin{cases} I_{it} + U_{it} & \text{if } a_{it} = 0 \text{ and } U_{it} \leq (C - I_{it}) \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

Equation 5 shows that if a customer is not visited (the first expression), the state, therefore the container status, will increase by the realized demand. If the customer is visited on a route or due to an emergency pickup (the second expression), the state for that customer in the next period will be zero, implying that the container is empty. Note that if a container is emptied in period t, then given the order of events, no new demand is realized after the container is emptied and therefore $I_{t+1} = 0$.

We modify the one period transition probability $\omega^{IOT}(I_{i(t+1)} \mid I_{it}, a_{it})$, the one period joint transition probability $\eta^{IOT}(\vec{I}_{t+1} \mid \vec{I}_t, d)$, and the one period loss function $L^{IOT}(\vec{I}_t, d)$ under the information policy as

$$\omega^{IOT}(I_{i(t+1)} \mid I_{it}, a_{it}) = \begin{cases} 1 & I_{i(t+1)} = 0 \text{ and } a_{it} = 1 \\ \phi_i(U_{it}) & I_{i(t+1)} = I_{it} + U_{it} \text{ and } a_{it} = 0 \\ 1 - \Phi_i(C - I_{it}) & I_{i(t+1)} = 0 \text{ and } a_{it} = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$\eta^{IOT}(\vec{I}_{t+1} \mid \vec{I}_t, d) = \prod_{i=1}^{N} \omega^{I}(I_{i(t+1)} \mid I_{it}, a_{it})$$

$$L^{IOT}(\vec{I}_t, d) = \underbrace{P\sum_{i=1}^{N}((1-a_{it})\cdot(1-\Phi_i(C-I_{it})))}_{\text{emergency pickup cost}} + \underbrace{p\sum_{i=1}^{N}((1-\Phi_i(C-I_{it})))}_{\text{loss of goodwill cost}} + \underbrace{h\sum_{i=1}^{N}\left((1-a_{it})\cdot \sum_{U_{it}<(C-I_{it})}(I_{it}+U_{it})\cdot\phi_i(U_{it})\right)}_{\text{holding cost}} \tag{6}$$

where the superscript I denotes the information policy hereafter. Note that the $T(d)$ does not change in the information policy. Then, the service provider's objective function in the information policy is given by

$$\pi^{IOT} = \min_{1\le d\le k} f^{IOT}(\vec{I}_t, d) + \bar{z} \tag{7}$$

where $f^{IOT}(\vec{I}_t, d) + \bar{z} = T(d) + L^{I}(\vec{I}_t, d) + \sum_{I_{t+1}} f^{I}(\vec{I}_{t+1}, d)\cdot \eta^{I}(\vec{I}_{t+1} \mid \vec{I}_t, d)$.

We summarize our main notation within a table located in Appendix A. For small-sized problems with up to six customers, an optimal solution to the MDP formulation can be calculated within a reasonable time (less than an hour). However, for larger problems, finding a solution becomes infeasible due to the curse of dimensionality inherent to dynamic programming. Nevertheless, the MDP serves two important purposes. It provides a structure and means of analysis to develop a heuristic policy of more practical relevance and it also enables a means for testing the heuristic's performance.

3.5. A Special Case with Single Customer

In this section, we define a special case to obtain a limited insight to the problem. The special case also provides a foundation to develop our heuristic and to derive some insights in assessing the value of the IOT application. Consider that the SDDSP serves only a single customer on a route. Moreover, as is typical in the secure data destruction industry, we assume that loss of goodwill cost p and holding cost h are negligible. Finally, we assume that the demand at the customer facility follows a Poisson distribution with mean λ. The objective function for the base policy can be expressed so as to minimize the average expected costs per visit cycle x, unless an emergency pickup occurs. We denote this policy as π_S^B where S denotes the special case and

$$\pi_S^B = \frac{2c_{01} + F + K + P\sum_{j=C+1}^{\infty}\phi^{x-1}(j)}{x}.$$

Note that average expected costs per visit cycle x are composed of the transportation cost, fixed truck setup cost, fixed collection cost, and the expected emergency pickup cost.

Proposition 1 *Unless an emergency pickup occurs, the optimal visit cycle x for the special case is given by* $x = \frac{2c_{01}+F+K+P\left[1-\Phi^{x-1}(C)\right]}{P\lambda\varphi^{x-1}(C)}$. *The corresponding optimal objective function is* $\pi_S^B = P\mu\phi^{x-1}(C)$.

Technical proofs are provided in Appendix B.1. Proposition 1 implies that, as expected, an increase in c_{01}, F, and K increases the optimal visit cycle. In contrast, by taking the derivative of x with respect to P, one can see that an increase in P decreases the optimal visit cycle.

Under the information policy, the special case problem is reduced to a periodic (s, S) system that has been shown to be optimal under different conditions (e.g. Iglehart (1963), Veinott and Wagner (1965), Beyer and Sethi (1999)). This system works in the following form: if the current inventory is higher than the threshold s, do not order; if lower, order such an amount to reach the target level S. Our implementation of this system differs from the literature in two ways. First, the SDDSP should visit the customer when the container status reaches s. Second, S is given by the capacity of the container C, therefore our task is to determine the optimal value of s.

Assume that a time period starts with container status $I_t \geq 0$. Then, following Veinott and Wagner (1965) but adjusting for emergency pickups, we define the expected emergency pickup cost incurred during a time period as

$$E_e(I, C) = P[1 - \Phi(C - I_t)]$$

Note that right after a visit to the customer, the next time period starts with $I_t = 0$. Now, let random variable $G(s)$ be the first period in which the cumulative demand exceeds s. Therefore, $G(s)$ represents the cycle time between two visits. Then, we can define the expected emergency pickup cost during $G(s)$ as

$$E(s, C) = E_e(0, C) + \sum_{i=1}^{\infty}\sum_{U_i=0}^{s} E_e(U_i, C)\phi^i(U_i)$$

where $E_e(U_i, C)$ is the expected emergency pickup cost in the $(i+1)$th period since the last visit, given that after $i > 0$ periods cumulative demand equals $U_i \leq s$ with probability $\phi^i(U_i)$. Simply, $E(s, C)$ denotes the convoluted effect of all possible realizations of s during $G(s)$.

The cumulative probability that over i periods demand does not exceed k equals $\Phi^i(k) = \sum_{t=0}^{k}\phi^i(t)$. The renewal function $M(s) = \sum_{i=1}^{\infty}\Phi^i(s)$ then gives the expected number of periods before cumulative demand exceeds s. Hence, the expectation of $G(s)$ becomes $1 + M(s)$. Now, we can write the objective function per period under the special case when container status information is available, denoted as π_S^{IOT}, as

$$\pi_S^{IOT} = \frac{2c_{01} + F + K + E(s, C)}{1 + M(s)}.$$

The optimal value of s, say s , minimizes this expression. However, we cannot find a closed-form expression for s . Nevertheless, we approximate the optimal objective function and present bounds on that approximation.

Proposition 2 *The approximated objective function $\tilde{\pi}_S^{IOT}$ is bounded with the following expression:*

$$\mu P\phi(C - s) \leq \tilde{\pi}_S^{IOT} \leq \mu P\phi(C - s + 1)$$

.

The closed-form expressions of the bounds for the approximated objective function can be used to limit the search space. Yet, they do not enable us to assess model sensitivity with respect to model parameters as the impact of a change in model parameters on the value of the term $\phi(C - s)$ is not clear. For instance, an increase in P will likely increase the system cost and decrease s . While the latter effect is likely to decrease the cost due to emergency pickups, it will also increase the visit frequency, leading to an increase in transportation cost. However, for a relatively large problem with multiple facilities, the service provider may visit a customer before or after its optimal container status threshold s to benefit from transportation cost savings. It is not clear whether or not these savings outweigh the cost increase. We can also make similar arguments for h and p when included in the model. Clearly, an optimal solution to the problem would consider the trade-off between (1) the transportation cost savings due to visiting multiple customers on the same route and (2) the cost increase due to visiting some customers at their suboptimal visit cycles. Therefore, to clarify our understanding of both policies, we now proceed to the next section to develop our heuristic that conceptually takes this trade-off into account.

4. Heuristic

As noted previously, optimal solutions to realistic problems will be left to future computing advances. Consequently, we must construct heuristics to compare the base and information policies. Unfortunately, existing heuristics for similar problems noted in the literature review are not sufficient for our purposes, and we must create heuristics for both policies. As suggested by referees, the details of heuristic construction are in Appendix C. A general sketch is provided here.

The major issue that makes a MDP solution infeasible are the sizes of both the decision space and the state space. In our heuristic, we focus on reducing these spaces. The decision space is reduced by only looking at a subset of routes that we believe would make business sense. Specifically, only routes that maximize the available shift time H and also maximize coverage of demand points are considered. This substantially reduces the number of routes that need be considered from the feasible set. This practical consideration reduces the decision set roughly 95

To begin, we compute the optimal solutions for all customers as single facility subproblems in which each facility (customer demand point) is considered in isolation of all others. Of course, because demand points have to be aggregated into routes, they cannot all be serviced at their "optimal" time. For example, demand points A and B may be next to each other, and while it would be optimal for demand point A to be picked up every 4 days and optimal for demand point B to be picked up every 6 days, that is impractical and very costly. Sending a truck on different schedules, one every four days and one every six days to pick up customers next to each other is cost prohibitive. However, if the customers are far apart, it may make sense to visit them on different intervals. If they are close together, some potential routes may involve sending a truck to each facility every four, five, or six days. For heuristic comparison we compute what we call a "relative cost" in Appendix C, a cost based on visiting that demand point at a different frequency than would be optimal in the myopic sense, viewed solely from the vantage point of that demand point in isolation. Incorporating the distance between customers, we calculate what we call a "transportation cost savings" from incorporating a demand point in a route versus visiting that demand point on a different route. The combination of these costs, summed over the demand points within the route, provides an estimate for how effective a route would be.

By its very nature, the heuristic policy evaluates the alternative decisions with respect to the collective state of all customers. Hence, both planned and unplanned (emergency pickup) visits will impact future heuristic decisions and therefore, (1) route decisions are not fixed and (2) visit cycles for customers may vary stochastically over time.

For interested readers, the specific heuristic algorithm is provided in Appendix C. Additionaly, Appendix D provides the evaluation of the heuristic policy performance for the base and information casesOverall, we find that the performance is excellent. We now proceed to assess the value of the IOT application in the next section.

5. Assessing the Value of IOT Application

The principal benefit of utilizing the container status information in decision making is a reduction in the uncertainty regarding container usage. In the base policy, the service provider must consider the tradeoff between visiting early and visiting late. If a customer is not serviced before the container fills, the service provider has to provide an emergency pickup service. To prevent emergency pickups, the service provider should choose a shorter visit cycle. However, shortening visit cycles will lead to higher transportation costs. As uncertainty increases, the tradeoff between these two costs becomes more pronounced. Generally, service providers tend to visit more frequently to help avoid emergency pickups.

The question remains, however, as to whether the reduction in uncertainty through the use of new technology that provides container status information translates into cost savings and what are the operating conditions that drive those cost savings, if any. Cost savings correspond directly to the value of the IOT application. We evaluate this value as the percentage improvement in expected cost achieved with information relative to the base scenario.

For the special case introduced in section 3.5, we can obtain the bounds for the value by combining π_S^B in Proposition 1 with the expression in Proposition 2. This gives us the following closed form expressions of the bounds on the value of the IOT application:

$$\mu P[\phi^{x\ -1}(C) - \phi(C - s\ )] \leq \pi_S^B - \tilde{\pi}_S^{IOT} \leq \mu P[\phi^{x\ -1}(C) - \phi(C - s\ + 1)]$$

These expressions enables us to only assess the sensitivity with respect to two parameters: the value of the IOT application increases with respect to μ and decreases with respect to C. The sensitivity with respect to other model parameters is not clear due to the term $[\phi^{x\ -1}(C) - \phi(C - s\ )]$. Hence, we clarify our understanding of the value of the IOT application through a simulation study.

5.1. Simulation Procedures

To evaluate the value of the IOT application, we consider a set of experiments that comprise a random design. We simulate each experiment with both the base policy and the information policy and measure the value of the IOT application as the cost difference between the two. We developed a custom simulation program using LabVIEW 2012 programming language to simulate the operational environment of the service provider under both base and information cases. Each experiment is simulated for 250 periods and replicated 40 times. The first 50 periods of each replication are set aside as the simulation warm-up period. Statistics are calculated for 200 periods in each replication. The warm-up period was chosen for convenience, yet larger than the number of days necessary for the system to exhibit steady-state behavior. Each period represents a day in the simulation. In each replication, the random number streams across all experiments are identical in order to reduce sampling error. We assess the accuracy of the simulation in measuring the value of the IOT in Appendix E.

5.2. Experimental Design

Given the sheer number of model parameters, we implement a randomized experimental design. The parameter values are partially based on values we observe in the data gathered by us from an SDDSP, but are also expanded to cover a broader range of values to reflect the range of operating cost environments that SDDSPs in general are likely to experience. Here, we discuss our method to generate parameters to evaluate the value of the IOT application. Specifically, we generate uniform random values for the following parameters: K $U(1,10)$, F $U(1,10)$, h $U(1,3)$, P $U(25,100)$, and p $U(5,25)$. The container capacity considered for the study is randomly chosen to be C $\{2,5,10\}$. We generate uniform random values for μ_i from $U(0.1,1.6)$, $U(0.25,4)$, and $U(0.5,8)$ for $C = 2$, $C = 5$, and $C = 10$, respectively. We set the fixed service time and the daily working hours across all experiments to $B = 0.25$ hours and $H = 8$ hours, respectively.

We also assess the value of the IOT application for different $\frac{\mu}{\sigma 2}$ using three different demand distributions: Poisson, Negative Binomial, and Binomial. For each scenario, we randomly select one of these three distributions. The maximum value for demand, n, when demand follows binomial distribution is also fixed at 15. The demand distributions used in the study are Pois(μ), B($n, \frac{\mu}{n}$), and NB($n - C, 1 - \frac{\mu}{n - C + \mu}$) for Poisson, Binomial, and Negative Binomial distributions respectively.

The number of customers considered for the study is chosen to be N $\{10,20,30,40,50\}$. For each value of N, we generated 10 different combinations for customer locations. Customer locations are randomly generated within a unit square. Thus, we generated 50 different combinations of customer locations. Transportation times and costs (b_{ij} and c_{ij}) are generated as explained in Appendix D. For each combination of customer locations, 100 random scenarios are generated which provide a total of 5,000 different scenarios. We simulate each scenario as explained in section 5.1 for both base and information cases and calculate the value of the IOT application using simulated average base case and information case costs. In the results we generate, the estimated standard error for the average period cost, using each heuristic, averages 1.4% of its mean value, and has a maximum error of 5.0%.

5.3. General Results and Observations

Overall, the value of the IOT application ranges between 0.0% and 20.8% with an average of 5.7%, median of 4.7%, and 10% of the experiments achieve a value that is more than 12.3%. In Figure 2 we break out all of the cost components using the averages across all experiments. The bar on the left (right) shows the total average cost per customer in the base (information) scenario. Each cost component is represented with a different pattern. The length of a bar for a cost component is equal to the average cost per customer for that component. On average, the total average cost per period per customer is $13.69 in the base scenario and $13.06 in the information scenario. Most of the contribution to the value comes from the savings in emergency pickup cost. On average, the information decreases emergency pickup costs by 19% (from $3.57 without information to $2.88 with information). That is, information increases the average time between successive emergency pickups per customer from 17.5 days to 21.7 days. From another perspective, the service provider can prevent one emergency pickup with information for every five emergency pickups without information. In essence, information enables better timing of visits so that emergency pickups are reduced. Note too that information only modestly reduces loss of goodwill by 3.8% which in total represent a smaller portion of average total cost than the other cost components. The reduction in loss of goodwill is much lower than the savings observed with respect to the emergency pickup cost. Since information enables more timely visits, by visiting when a container is nearing capacity, a container may fill when a visit is already scheduled.

Figure 2 Average cost breakdown

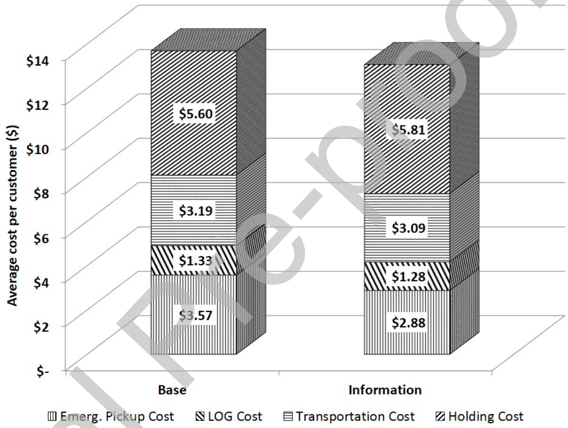

The flip side of enabling more timely visits to customers is that information helps to reduce transportation cost by avoiding unnecessary visits to customers. Unnecessary visits occur when a visit is made to empty a container that is already near empty. With information the service provider can avoid these unnecessary visits, even if quite some time has elapsed since the last visit. Transportation costs are on average 3.1% lower (from $3.19 without information to $3.09 with information). Clearly, as the number of unnecessary visits decreases, so too will the corresponding transportation cost. However, a by product of visiting less frequently is that documents are held for a longer period of time. Consequently, we observe a modest increase in holding cost. As is shown in Figure 2, the information increases the holding cost per period, on average, 3.8% (from $5.60 without information to $5.81 with information).

While the average cost analysis we provide here is representative of the results we observe across experiments, it is clear that the value of the IOT application is sensitive to model parameters as indicated by the range of the value demonstrated in our results. We now proceed with a sensitivity analysis.

5.4. Sensitivity Analysis

In this section, we report a sensitivity analysis on the value of the IOT application with respect to model parameters. Table 1 summarizes our results. The intervals for the parameter values are given in Table 3 located in the appendix. The reported values are averages across all experiments that contain a value for a parameter indicated by the row header and lies within the interval indicated by the column header. For example, on average, the value of the IOT application is 5.7% when the holding cost is medium (between \$1.65 and \$2.35). Similarly, when a truck serves 30 customers, the average value is 6.0%. Note that given our long simulation runs, with 40 replications using common random number streams, all values reported in Table 1 are statistically different from zero, with a maximum p-value < 0.01 using paired t-tests.

Table 1 Sensitivity analysis for the value of the IOT application

	Parameter interval				
Parameter	Low	Medium	High		
p	5.7%	5.6%	5.7%		
P	4.5%	5.6%	7.0%		
h	6.9%	5.7%	4.6%		
F	5.7%	5.8%	5.7%		
K	6.0%	5.7%	5.5%		
c_{ij}	4.8%	4.6%	4.5%		
$\frac{\mu}{\sigma^2}$	6.6%	5.6%	5.1%		
C	10.6%	4.7%	1.8%		
	Number of customers				
	10	20	30	40	50
N	4.6%	4.9%	6.0%	6.5%	6.5%

The results reported in Table 1 largely confirm our intuition. Average value increases with respect to P and decreases with respect to h, $\frac{\mu}{\sigma^2}$, and C (differences between levels for each of these parameters are statistically different at p-value < 0.01 using paired t-tests). The value of the IOT application is not sensitive to p, c_{ij}, K, and F (p-values> 0.05).

It is somewhat surprising that the value of the IOT is largely insensitive with respect to the transportation cost (c_{ij}). One would expect that as transportation cost increases, with information, there would be more opportunity to save, since the service provider would decrease the number of unnecessary visits. However, we do not observe this behavior in the results. If anything, we observe the reverse. When it is expensive to visit customers, the service provider prefers denser routes to leverage transportation pooling so that the average transportation cost per customer is low. The cost savings provided through information by enabling more timely visits may not be as significant as that provided through transportation pooling. Hence, it may be more beneficial to include a customer on a route, even though it may not be *timely* to visit. In these situations, information that identifies a container is not near capacity is ignored (customer is visited regardless) and consequently the value is less.

The value of information increases with respect to the emergency pickup cost P. This is not surprising given our general observations regarding average cost performance across experiments. Since the largest contribution to the value arises from reducing emergency pickups, an increase in P generates greater value (from \$0.39 savings per customer when P is low to \$0.96 savings per customer when P is high).

Small $\frac{\mu}{\sigma^2}$ corresponds to a planning environment with high uncertainty. At the extreme, if the system is completely deterministic with adequate capacity there would be no value as all visits could be timed perfectly without information and there would be no emergency pickups or loss of goodwill through poor service. It is intuitive therefore as uncertainty increases, costs increase, as well as the opportunity to reduce them through information.

5.5. Route Capacity and Geographical Dispersion

The analysis up to this point has been done using random customer locations and a fixed routing capacity in hours. However, the value may also depend on different routing characteristics such as the geographical dispersion of customers and route capacity. Geographical dispersion is associated with the relative locations of the customers. When the customers are close to each other and to the service provider, we can say that the dispersion is low. Therefore, transportation time and cost is lower and more facilities can be visited in a period than when they are further away from each other and the service provider. When dispersion is high, we observe longer traveling times between each facility.

Route capacity is associated with the available time to visit customers and by extension, the number of customers that can fit on a route. In our model, route completion time is restricted with the total number of working hours H. Every additional working hour may allow visiting more customers in the same route. Does the information become more valuable when customers are dispersed? How much value can be obtained if the working hours are increased? To answer these questions, we conduct an additional numerical study.

To evaluate the impact of dispersion and capacity on the value of the IOT application, we conduct a 2x2 factorial design experiment to assess the impact of these two routing characteristics on the value at two different levels of each characteristic: low and high. Hence we have four cases: high dispersion-low capacity, low dispersion-low capacity, high dispersion-high capacity, and low dispersion-high capacity. The case generated with the assumptions that are described below is the "high dispersion-low capacity" case. To look at the "low dispersion-low capacity" case, we multiply the distances of all customers from the service provider (c_{0i}) generated in the high dispersion-low capacity case by 0.9

and recalculate b_{ij}, c_{ij}, and the route decision space. To look at the "high dispersion-high capacity" case, we increase the total working hours H used in the high dispersion-low capacity case by one hour. To look at the "low dispersion-high capacity" case, we change both H and c_{0i} as explained above. For each case, we keep all other parameters the same as in the high dispersion-low capacity case. Moreover, we replicate the 2x2 experimental design for different values of N, since the number of customers will directly impact both dispersion and the use of capacity.

For the numerical study, we choose the number of customers to be $N \in \{10, 20, 30\}$. For each value of N, we generate 100 different scenarios for customer locations. Thus, we generated 300 different scenarios with different combinations of customer locations. For model parameters, in all cases, we use the middle value of each parameter range defined in section 5.2. Transportation times and costs (b_{ij} and c_{ij}) are generated as explained in Appendix D. We use the Negative Binomial distribution to generate mean demands. The total number of working hours H in the high dispersion-low capacity case is set to 8 hours.

We simulate each scenario as explained in section 5.1 for the four different cases of geographic dispersion and routing capacity. Table 2 summarizes the results for the value of the IOT. For instance, the average value with 20 customers when the dispersion is low and the route capacity is high is 1.4%.

Table 2 The value of IOT with dispersion and route capacity

	Capacity (N = 10)		Capacity (N = 20)		Capacity (N = 30)	
Dispersion	Low	High	Low	High	Low	High
Low	1.7%	1.3%	1.7%	1.4%	1.9%	1.6%
High	2.2%	1.5%	2.4%	1.6%	2.8%	1.8%

Table 2 shows that when there is low geographical dispersion, the value is lower. When customers are close to each other, due to transportation pooling, it becomes more profitable to add a customer on a route even though it may be early to visit (container is not filled to near capacity). This creates higher route density. In this case, the savings from the transportation cost by adding a customer to the route are more than the savings that can be obtained by the information mainly by visiting that customer in a later time period. Therefore, there is less value with low dispersion. When dispersion is high, customers are located far from each other and it is not possible to visit multiple customers in a single route. Therefore, the information obtained from the containers becomes more valuable.

According to Table 2, high route capacity leads to low value. If the maximum working hours per period increases, it becomes possible to add more customers to a route and hence, routes become more dense. As in the case of low geographical dispersion case, this leads to lower value. When there is both high route capacity and low dispersion, it is not surprising that this is the case in which the value is least. When the number of customers increases, the impact of low dispersion and high route capacity decreases and the value increases.

6. Conclusion

This study evaluates the effectiveness of a new IOT application that had the potential to strategically change a business model in the secure data destruction industry. The executives supporting the IOT invention at this firm believed that this would be truly revolutionary for their business. Without such technology, the container status information is only available after a truck and driver are dispatched and on-site. In that case, the routing policy depends on historical demand rates and how long it has been since the last visit. Because of demand uncertainty, routing decisions will invariably entail visiting customers earlier or later than the best time to visit. Visiting early results in increased transportation costs and visiting late results in poor service. By virtually eliminating demand uncertainty, smart containers enable dynamic routing to customers on an as needed basis. As a general statement, more information is generally better than less. However, the specific IOT application here got its information at too high a cost. Further, the nature of dynamic routing would have necessitated the firm change considerably. Ultimately, this IOT application did not occur.

However, commercial success is not the sole hallmark of research progress. Here, we have studied an interesting system in which (1) stochastic demand occurs in containers at remote locations, (2) customer inventory levels (i.e. container fill levels) are known prior to route decision making only if the new IOT application is implemented, otherwise they are not known, and (3) emergency pickup occurs when demand exceeds capacity. Problems with these three characteristics have not previously been addressed in the literature. We assess the value of the new IOT application under different parameter settings and routing conditions. The benefits of the IOT application ranges upwards to as high as 20.8%. The greatest value is derived from reducing the occurrence of emergency pickups. We observe an average reduction of 19% in emergency pickups in our simulation study. Further, information decreases transportation costs by 3.1% and loss of goodwill cost by 3.8% at the expense of a 3.8% increase in holding cost. We also show that the value decreases when there is low dispersion and high route capacity.

We demonstrate that the value of the new IOT application is not straightforward and depends on certain parameter choices. To make an investment for the smart container technology, SDDSPs should be careful about the values of each parameter in their operating environments. The value will be higher if (1) the emergency pickup cost is high, (2) the holding cost is low, (3) the demand uncertainty is high, (4) the daily working hours are not flexible, and (5) customers are geographically dispersed. Under these conditions, SDDSPs may expect higher return on the technology investment.

Emergency pickup/delivery is common in practice. Even though we develop our model within the context of the SDDSP industry, our work can readily be adapted to address similar problems in different industries including propane/oil delivery, medical waste collection, recycling, waste management, blood product delivery, etc. There are similar technologies in these industries to that of smart containers in the SDDSP industry.

There are several important avenues for future research that extend to include modeling multiple products. For example, SDDSPs also collect old hard drives from their customers in addition to sensitive documents. Further, they generally have multiple vehicles serving different regions. Another promising avenue for future research is the level of the sensor technology used in the containers. In our model, we assume that sensors are able to detect every 10% change which corresponds to one of the current technologies available. However, in practice, sensors have different levels of resolution (ranges between 0.3% and 75%) and may provide data with higher resolution or lower resolution depending on the technology used. We conjecture that high resolution sensors provide more value.

Appendix A: Table of Parameters

Symbol	Notation	Symbol	Notation
N	number of customers	C	secured container capacity
K	fixed collection cost	F	truck setup cost
h	holding cost	d_t	route decision at time t
H	maximum working hours per period	s_d	number of customers on route d
μ_i	mean demand at customer i	x_{it}/I_{it}	state level for customer i at time t
p	loss of goodwill cost	P	emergency pickup cost
τ_i	SFVC/SFVP for customer i	$q_{i\tau_i}$	cost difference between the costs for τ_i and that for τ_i

Appendix B: Proofs of Propositions

B.1. Proposition 1

We can rewrite the objective function as

$$\pi_S^B = \frac{2c_{01} + F + K + P[1 - \Phi^{x-1}(C)]}{x} \tag{8}$$

From Poisson distribution, we have $\phi^x(y) = \frac{e^{-\mu x}(\mu x)^y}{y!}$ and $\Phi^x(z) = \sum_{y=0}^{z}\left[\frac{e^{-\mu x}(\mu x)^y}{y!}\right]$. Note that technically x is an integer, but here we assume a continuous approximation. The derivative of $\Phi^x(z)$ with respect to x is

$$\frac{\partial \Phi^x(z)}{\partial x} = \sum_{y=0}^{z}\left[\frac{-\mu e^{-\mu x}(\mu x)^y}{y!} + \frac{e^{-\mu x}\mu y(\mu x)^{y-1}}{y!}\right] = -\mu\sum_{y=0}^{z}\left[\frac{e^{-\mu x}(\mu x)^y}{y!}\right] - \mu\sum_{y=0}^{z-1}\left[\frac{e^{-\mu y}(\mu y)^y}{y!}\right] = \frac{-\mu e^{-\mu x}(\mu x)^z}{z!} = -\mu\phi^x(z)$$

Therefore,

$$\frac{\partial \pi_S^B}{\partial x} = -\frac{2c_{01} + F + K + P}{x^2} + \frac{P\left[x\mu\phi^{x-1}(C) + \Phi^{x-1}(C)\right]}{x^2}$$

It follows that

$$\frac{\partial \pi_S^B}{\partial x} = 0 \Leftrightarrow x = \frac{2c_{01} + F + K + P\left[1 - \Phi^{x-1}(C)\right]}{P\mu\phi^{x-1}(C)}$$

Plugging x in Equation 8 reveals that $\pi_S^B = P\mu\phi^{x-1}(C)$. □

B.2. Proposition 2

Suppose that the starting container status in the last period of a cycle equals exactly s. Then, the expected emergency pickup cost incurred in a cycle in the last period becomes $E(s,C) \approx E_e(s,C)$. The corresponding cycle time is approximately $1 + s/\mu$. Hence the approximation for average profit per period when container status information is available is

$$\pi_S^{IOT} \approx \tilde{\pi}_S^{IOT} = \frac{2c_{01} + F + K + E_e(s,C)}{1 + s/\mu} \tag{9}$$

If s is optimal, then naturally $\tilde{\pi}_S^{IOT}(s) \le \tilde{\pi}_S^{IOT}(s - 1)$ and $\tilde{\pi}_S^{IOT}(s) \le \tilde{\pi}_S^{IOT}(s + 1)$. Plugging Equation 9 into each inequalities gives the following expression for the upper and lower bounds of $\tilde{\pi}_S^{IOT}$:

$$\mu P\phi(C - s) \le \tilde{\pi}_S^{IOT} \le \mu P\phi(C - s + 1) \tag{10}$$

□

Appendix C: Algorithm for the Heuristic Policy

C.1. Reducing Decision Space

Due to economies of scale, optimal solutions to IRPs tend to include dense routes. Dense routes are the routes that are near their capacity. These routes generate considerable savings in transportation cost. Therefore, we reduce the decision set to include only the most dense routes. To obtain the reduced decision set, we first group routes based on the number of customers. We call each group a *level* and all the routes in a given level contain the same number of customers. Thus, the highest level includes routes with the highest density. Only routes in which all customers can be serviced within H hours are considered. Then, we identify the first level below the highest level in which all facilities are included in at least one route at that level. We refer to this level as *the maximum feasible level*. The reduced decision set includes all routes at the maximum feasible level and in any level above that. The reduced decision set also includes the empty route.

Our approach reduces the decision space significantly. Consider a random example with 30 customers. Figure 3 shows the total number of decisions for each level. In this example, the maximum feasible level is 7 and the highest level is 9. Thus, the reduced decision set includes only 365=2+49+313+1 routes, which is approximately 4% of the number of routes in the original feasible decision set.

C.2. Myopic Policy

The heuristic policy is myopic since it only considers the current period. It is predicated on the optimal solutions of single facility subproblems in which each facility is considered in isolation. Dense routes in the reduced decision set will decrease overall transportation cost. However, they may unnecessarily impose visits to some customers more or less frequently than needed, resulting in an increase in average cost per period. Our heuristic evaluates each route in the reduced decision set based on (1) the transportation cost savings that arise from visiting multiple customers in the same route relative to visiting each customer individually and (2) the cost increase due to visiting some customers earlier or later than the optimal planned visit cycle. We refer to the latter cost as the relative cost (RC). We discuss each below, beginning with the transportation cost savings.

Figure 3 A sample decision space

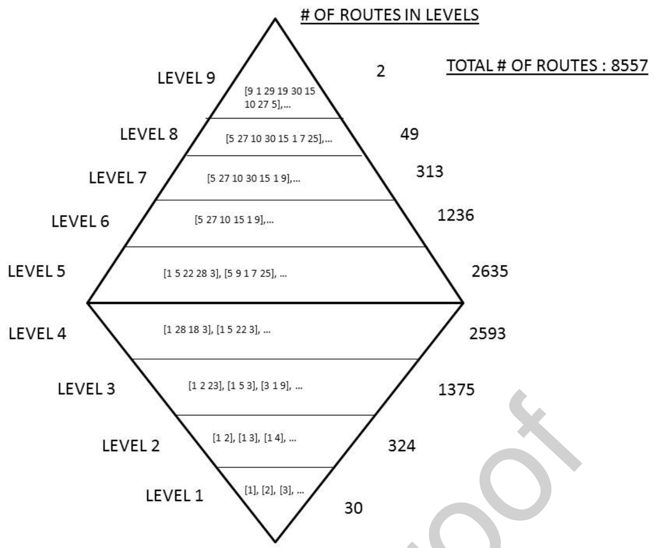

C.2.1. Transportation Cost Savings The transportation cost savings for a decision arise when there is more than one customer in the same route. The savings come from the shared truck setup cost F and, potentially, a decrease in total traveling cost $\sum c_{ij}$. Therefore, the transportation cost savings for route d is given by

$$TCS(d) = \begin{cases} \underbrace{(s_d - 1) \cdot F}_{\text{setup savings}} + \underbrace{(2\sum_{j=1}^{s_d} c_{0j}) - TC_d}_{\text{transportation cost savings}} & \text{if } s_d > 1 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

The first expression to the right of Equation 11 defines the shared truck setup cost. Each additional customer on the route increases savings by F. The second expression provides the total transportation cost if each customer is visited individually. The third expression defines the total traveling cost for route d. Hence, the difference between the second and the third expression gives the potential decrease in total traveling cost when visiting multiple customers in the same route relative to visiting them in isolation.

C.2.2. Relative Cost To calculate the relative cost in the base scenario, we first decompose the IRP into single facility subproblems. For each subproblem, we solve the MDP formulation in Equation 4 and the solution provides an optimal single facility visit cycle (SFVC). An optimal SFVC, denoted as τ_i for customer i, is defined as the first countable state starting with state 1 in which the decision is to visit. For example, if an optimal policy for a single facility problem is $\{0, 0, 0, 1, 1, \ldots\}$ (where 0 indicates not to visit and 1 indicates to visit) for the states $\{1, 2, 3, 4, 5, \ldots\}$, then $\tau_i = 4$. This means that unless an emergency pickup occurs, the service provider will visit the customer every four periods. Let $\vec{\tau} = [\tau_1, \tau_2, \ldots, \tau_N]$ denote optimal SFVCs for all N customers.

While we use Equation 4 to obtain τ_i, we can constrain the MDP to obtain the average cost per period for any SFVC. In this manner, we can assess the cost of visiting earlier or later than the optimal SFVC. For instance, if we want to obtain the average cost per period for a SFVC $= 2$ for the example above, we can solve the MDP by imposing the policy $\{0, 1, 1, 1, 1, \ldots\}$. Now, for customer i, let τ_i be any SFVC. Moreover, let $q_{i\tau_i}$ denote the increase in average cost per period in the single facility subproblem when the SFVC is τ_i relative to that when

SFVC is τ_i . Hence, $\vec{q_i} = [q_{i1}, q_{i2}, \ldots]$ contains information about the incremental cost of visiting early or late for customer i when the SFVC is different from the optimal SFVC. Thus, $q_{i\tau_i} = 0$.

The RC^B is a function of $\vec{x_t}$ and d. The relative cost for a given route includes (1) costs due to early/late visits to all customers that are on the route and (2) benefits from delaying the visit to all customers that are not on the route. Given that the state is $\vec{x_t}$, the visit cost for customer i is simply equal to $q_{ix_{it}}$. The benefit from delaying the visit to customer i is given by $q^D_{ix_{it}} \equiv q_{ix_{it}} - q_{ix_{i(t+1)}}$ and it conceptually represents the change in average cost per period by delaying the visit decision one period in state x_{it}. Note that this benefit can be negative, implying that delaying the visit decision to a customer may increase the RC^B for that route. For ease of exposition, let $q_{ix} \equiv q_{ix_{it}}$ and $q^D_{ix} \equiv q^D_{ix_{it}}$. Finally, let A be an arbitrarily large number. Then, $RC^B(\vec{x_t}, d)$ is given by

$$RC^B(\vec{x_t}, d) = \sum_{i=1}^{N} \begin{cases} \underbrace{(a_{it} \cdot q_{ix})}_{\text{visit cost}} - \underbrace{(1 - a_{it}) \cdot q^D_{ix}}_{\text{delay benefit}} & \text{if } x_{it} < \tau_i \\ \underbrace{-a_{it} \cdot [q_{ix} + A \cdot (x_{it} + 1 - \tau_i)]}_{\text{visit cost}} - \underbrace{(1 - a_{it}) \cdot \left[q^D_{ix} - A \cdot (x_{it} + 1 - \tau_i)\right]}_{\text{delay benefit}} & \text{if } x_{it} \geq \tau_i \end{cases} \quad (12)$$

Costs arise from both visiting and not visiting customers and are largely predicated on the difference between $\vec{x_t}$ and $\vec{\tau}$. The first expression to the right of Equation 12 concerns any customer whose current state x_{it} is before its optimal SFVC. This expression defines the additional cost of visiting a customer early (first term) or the savings that arises from delaying the visit decision one period towards its optimal SFVC (second term). Conceptually, this expression aims to prevent early visits.

The second expression of the equation concerns any customer whose current state x_{it} is at or beyond its optimal SFVC. Conceptually, this expression aims to prevent late visits. It does so by ensuring that any further decision to delay visiting a late customer will increase the priority to visit that customer in the future. Note that if $x_{it} = \tau_i$ and the decision is to visit, RC^B reduces to $-A$. For each additional period of delay beyond τ_i , the decision to visit a late customer will decrease RC^B by $(A + q_{ix})$, making the routes that include late customers more attractive. Conversely, if the decision is not to visit (the second term in the expression), then RC^B increases by $(A - q^D_{ix})$ for each additional period beyond τ_i , making the route unattractive. Some additional discussion is warranted here. By definition, $q_{ix_{it}}$ includes the additional cost for visiting customer i at its non-optimal SFVC. Note that for $x_{it} > \tau_i$, the additional cost tends to flatten and approaches a constant as x_{it} increases. This arises because the probability of emergency pickup approaches one. Hence, using an arbitrarily large number A will ensure that routes visiting customers near their optimal SFVC are attractive.

For the information policy, the RC^I is a function of $\vec{I_t}$ and d. The only difference in calculation RC^I is that we solve Equation 7 for each customer to obtain an optimal single facility visit point (SFVP). Therefore, τ_i represents SFVP. Similar to an optimal reorder point in the (s, S) system, an optimal visit point triggers the visit decision to a customer. For example, if an optimal policy for a single facility problem in the information case is {0, 0, 0, 1, 1, . . .} (where 0 indicates not to visit and 1 indicates to visit) for the states {0, 1, 2, 3, 4, . . . C}, then τ_i is 3 since a visit is triggered once the container status reaches the value 3. This means, unless an emergency pickup occurs, the service provider will visit the customer whenever the container fill level reaches three. The rest of the calculation for RC^I is identical RC^B . Note that the benefit from delaying the visit to customer i in the information policy conceptually represents the change in average cost per period by delaying the visit decision to the next container status. We modify $RC^I(\vec{I}_t, d)$ in the information policy as

$$RC^I(\vec{I}_t, d) = \sum_{i=1}^{N} \begin{cases} \underbrace{(a_{it} \cdot q_{il})}_{\text{visit cost}} - \underbrace{(1 - a_{it}) \cdot q_{il}^D}_{\text{delay benefit}} & \text{if } I_{it} < \tau_i \\ \underbrace{-a_{it} \cdot (q_{il} + A \cdot (I_{it} + 1 - \tau_i))}_{\text{visit cost}} - \underbrace{(1 - a_{it}) \cdot \left(q_{il}^D - A \cdot (I_{it} + 1 - \tau_i)\right)}_{\text{delay benefit}} & \text{if } I_{it} \geq \tau_i \end{cases} \tag{13}$$

The heuristic decision in the base (information) policy for state $\vec{x}_t$ ($\vec{I}_t$) is simply the route d that minimizes the di□erence between the RC^B (RC^I) and the $TCS(d)$. Therefore, the heuristic objective function for a state $\vec{x}_t$ ($\vec{I}_t$) is given by π^{heu} where the superscript *heu* denotes the heuristic, denotes the reduced decision set, and

$$\pi^{heu} = \min_{d} \{RC - TCS(d)\} \tag{14}$$

C.2.3. Heuristic Description To find the heuristic decision for a given state, we (1) generate the reduced decision set, (2) generate optimal SFVC/SFVP vector ($\vec{\tau}$) and incremental cost vectors ($\vec{q}_i$), and (3) compare routes based on Equation 14, respectively.

- *PROCEDURE 1: Generate the reduced decision set*

  STEP 1. Generate all feasible routes

  STEP 2. Group routes based on the number of customers (i.e. generate levels).

  STEP 3. Identify the highest level in which all facilities are included in at least one route (i.e. the maximum feasible level).

  STEP 4. In addition to the empty route, the reduced decision set includes all routes in the maximum feasible level and in any level above that.

- *PROCEDURE 2: Generate $\vec{\tau}$ and $\vec{q}_i$*

  STEP 1. For each customer, define a single facility subproblem.

  STEP 2. For each subproblem, solve Equations 4 and 7 for optimality for the base and the information policies and identify the first countable state in which the decision is to visit (i.e. τ_i) for each policy. This step will generate $\vec{\tau}$.

  STEP 3. For each subproblem, solve Equations 4 and 7 by imposing policies associated with di□erent SFVCs/SFVPs (i.e. τ_is). For instance, $\tau_i = 1$ corresponds to the policy of $\{0, 1, 1, 1, 1, \ldots\}$, $\tau_i = 2$ corresponds to the policy of $\{0, 0, 1, 1, 1, \ldots\}$, and so on.

  STEP 4. For each customer i, compare the average costs obtained in step 3 to the optimal average costs per period to generate $\vec{q}_i$ for both policies.

- *PROCEDURE 3: Evaluate routes*

  STEP 1. For each route, calculate the value of Equation 14 using Equation 11 and the relevant RC.

  STEP 2. The heuristic decision for a policy is the route that has the lowest $RC - TCS(d)$ value.

Appendix D: Heuristic Performance

We evaluate the heuristic policy by calculating Equation 14 for the base and the information polices separately and comparing the heuristic costs to the optimal costs specified in Equations 4 and 7, respectively. We consider a set of experiments that comprise a random design. In our experimental design, periods are days.

We use three different demand distributions: Poisson, Negative Binomial, and Binomial. Numerous studies on inventory modeling assume that the stochastic arrival of demand follows Poisson processes and thus the count of demand conforms to Poisson distribution. However, the Poisson distribution assumes equi-dispersion that makes it highly restrictive in empirical modeling if data exhibits the commonly-seen over-dispersion or under-dispersion. Therefore, in addition to the Poisson distribution, we also test the heuristic using Binomial and Negative Binomial distributions to represent the demand with low uncertainty and high uncertainty, respectively. Further, the Negative Binomial distribution characterizes retail demand sufficiently well (Agrawal and Smith 1996).

We generate uniform random values for the following parameters: $\mu \sim U(0.5, 8), K \sim U(1, 10), F \sim U(1, 10), h \sim U(1, 3), P \sim U(25, 100)$, and $p \sim U(5, 25)$, where μ denotes the mean for any distribution. The capacity of the secured containers $C = 10$, the fixed service time $B = 0.25$ hours, and daily working hours $H = 8$ hours across all experiments. The maximum value for demand, n, when demand follows binomial distribution is also fixed at 15. The demand distributions used in the study are Pois(μ), B($n, \frac{\mu}{n}$), and NB($n - C, 1 - \frac{\mu}{n - C + \mu}$) for Poisson, Binomial, and Negative Binomial distributions respectively. We use data from the SDDSP to determine the values of certain model parameters such as K, F, h, and P. For other model parameters, we choose parameter values to test the robustness of the heuristic over widely varying operating cost environments including the operating environment of our SDDSP.

We generate random customer locations within a square, with the service provider located in the center of this square. The length of the edge of the square is chosen such that the time to travel between two customers that are located at two adjacent corners of the square is 5.48 hours. This allows the service provider to visit and serve any customer in the square within a period. For b_{ij} values, we first generate uniform random cartesian coordinates for $b_{0j} \sim U(-2.74, 2.74)$, then calculate the transportation time between each customer. The transportation cost to travel between facilities i and j, c_{ij}, is proportional to the time to travel between customers (b_{ij}). To assess the impact of different transportation costs on heuristic performance, we obtain uniform random values for c_{ij} by multiplying b_{ij} with a random proportion generated from $U(3, 9)$.

While the selected ranges of parameter values represent widely varying operating cost environments, our ability to solve the MDP for optimality is constrained by the number of customers which has been our primary motivation for developing our heuristic in the first place. The number of customers considered for the study is chosen to be $N \in \{2, 3, \ldots, 6\}$. For each value of N and each demand distribution, 200 random scenarios are generated which provide a total of $3,000$ scenarios. We use value iteration to compute the average expected cost for the MDP. For the heuristic, we solve the decision for each possible state and then use the state transition probabilities to solve for the long-run average cost.

We measure heuristic performance as the percentage increase in cost relative to the optimal cost. Overall, the heuristic for the base policy (information policy) achieves, on average, an expected cost that is 1.4% (1.3%)

greater than optimal, a worst case cost that is 11.7% (10%) greater than optimal, and 90% of the experiments achieve a cost that is no more than 3.2% (3.7%) greater than optimal.

We performed a sensitivity analysis for heuristic performance with respect to model parameters. Since the experiments comprise a random design, we categorize parameters as low, medium, and high. The corresponding intervals for each parameter are given in Table 3.

Table 3 Parameter categories

Parameters	Parameter intervals		
	Low	Medium	High
p	[5,11.5]	(11.5, 18.5]	(18.5, 25]
P	[25,50]	(50, 75]	(75, 100]
h	[1,1.65]	(1.65, 2.35]	(2.35, 3]
F	[1,4]	(4, 7]	(7, 10]
K	[1,4]	(4, 7]	(7, 10]
c_{ij}	[3,5]	(5, 7]	(7, 9]
$\frac{\mu}{\sigma^2}$	<1 (Neg. Binom.)	= 1 (Poisson)	>1 (Binomial)

Table 4 presents sensitivity analysis of the heuristic performance under both policies with respect to model parameters. For example, the heuristic in the information case achieves, on average, an expected cost that is 1.4% greater than optimal when the emergency pickup cost is between $50 and $75 (medium). As can be seen in these tables, the heuristic performance increases with respect to increases in p and h, decreases with respect to increase in c_{ij}, and does not appear to be sensitive to K, F, P, N, and $\frac{\mu}{\sigma^2}$. Overall, the performance is excellent.

Table 4 Sensitivity analysis (Heuristic vs. optimality)

Parameters	Base			Information		
	Low	Medium	High	Low	Medium	High
p	1.5%	1.3%	1.2%	1.7%	1.4%	1.2%
P	1.1%	1.3%	1.6%	1.4%	1.4%	1.4%
h	1.5%	1.3%	1.3%	1.8%	1.4%	1.1%
F	1.3%	1.3%	1.3%	1.4%	1.4%	1.5%
K	1.3%	1.3%	1.4%	1.4%	1.4%	1.5%
c_{ij}	1.2%	1.4%	1.4%	1.1%	1.4%	1.8%
$\frac{\mu}{\sigma^2}$	1.2%	1.5%	1.3%	1.3%	1.5%	1.5%
Policy	Number of customers					
	2	3	4	5	6	
Base	1.2%	1.2%	1.5%	1.4%	1.4%	
Information	1.5%	1.3%	1.3%	1.4%	1.6%	

Appendix E: Accuracy of the Simulation in Assessing the Value of IOT

To test the accuracy of the simulation in measuring the value of IOT, we duplicate the same 3000 scenarios we used to test heuristic performance in Appendix D. We obtain three different measures for the value for each scenario: optimal, heuristic, and simulation. To obtain optimal value, we computed the average expected cost for the MDP using value iteration under both base and information cases. For the heuristic value, we solve the heuristic decision for each possible state and then, use the state transition probabilities to solve for the long-run average cost for both base and information cases. Finally, for each scenario, we simulated both cases with random number streams using heuristic policies.

Table 5 demonstrates the value between (1) optimal, (2) heuristic, and (3) simulation across percentiles of the test cases which are ordered from the lowest to the highest value. For example, at the 50th (median) percentile, the heuristic value under Poisson demand is 3.3%, the simulation value is 1.8%, and the corresponding optimal value is 3.2%. Across all percentiles, the heuristic performance is ordered very closely to optimal performance. Note too that the value computed from the simulation is noticeably smaller across all experiments. We are unsure why this is the case. Nevertheless, the smaller value obtained through the simulation tempers our analysis, providing a conservative assessment of the results.

Table 5 The value of IOT for small-sized problems

	Neg. Binom. ($\frac{\mu}{\sigma^2} < 1$)			Poisson ($\frac{\mu}{\sigma^2} = 1$)			Binomial ($\frac{\mu}{\sigma^2} > 1$)		
Percentile	Opt.	Heur.	Sim.	Opt.	Heur.	Sim.	Opt.	Heur.	Sim.
0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
5	0.0%	0.0%	0.0%	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%
10	1.0%	0.6%	0.0%	0.4%	0.2%	0.0%	0.1%	0.0%	0.0%
25	2.3%	2.1%	0.9%	1.6%	1.2%	0.4%	0.6%	0.3%	0.0%
50	4.9%	4.8%	2.6%	3.2%	3.3%	1.8%	2.3%	2.2%	1.9%
75	8.1%	8.2%	4.6%	5.7%	5.7%	3.4%	4.7%	4.9%	3.1%
90	10.9%	10.8%	6.9%	7.4%	7.9%	5.3%	7.5%	7.8%	5.3%
95	13.1%	13.4%	8.9%	8.8%	9.3%	7.0%	9.1%	9.5%	6.7%
100	25.1%	25.3%	20.7%	15.4%	18.3%	13.6%	18.6%	18.1%	17.1%

References

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

Andersson, H., A. Ho, M. Christiansen, G. Hasle, A. Lokketangen. 2010. Industrial aspects of literature survey: Combined inventory management and routing. *Computers & Operations Research* **37** 1515–1536.

Atzori, L., A. Iera, G. Morabito. 2010. The Internet of Things: A survey. *Computers Networks* **54**(15) 2787–2805.

Aviv, Y. 2002. Gaining benefits from joint forecasting and replenishment processes: The case of auto-correlated demand. *Manufacturing & Service Operations Management* **4**(1) 55–74.

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

Beyer, D., S.P. Sethi. 1999. The classical average-cost inventory models if iglehart and veiott-wagner revisited. *J. Optim. Theory Appl.* **101**(3) 523–555.

Birpinar, M.E., M.S. Bilgili, T. Erdogan. 2009. Medical waste management in turkey: A case study in istanbul. *Waste Management* **29** 445–448.

Chen, L., H.L. Lee. 2009. Information sharing and order variability control under a generalized demand model. *Management Science* **55**(5) 781–797.

Chen, Y.M., C.T. Lin. 2009. A coordinated approach to hedge the risks in stochastic inventory-routing problem. *Computers & Industrial Engineering* **56** 1095–1112.

Cheung, K.L., H.L. Lee. 2002. The inventory benefit of shipment coordination and stock rebalancing in a supply chain. *Management Science* **48** 300–306.

Coelho, L.C., J-F. Cordeau, G. Laporte. 2014. Thirty years of inventory routing. *Transportation Science* **48**(1) 1–19.

Dror, M., M. Ball. 1987. Inventory/routing: Reduction from an annual to a short-period problem. *Naval Research Logistics* **34** 891–905.

Dror, M., M. Ball, B. Golden. 1985. A computational comparison of algorithms for the inventory routing problem. *Annals of Operations Research* **4** 3–23.

Faccio, M., A. Persona, G. Zanin. 2011. Waste collection multi objective model with real time traceability data. *Waste Management* **31** 2391–2405.

Federgruen, A., P. Zipkin. 1984. A combined vehicle routing and inventory allocation problem. *Operations Research* **32**(5) 1019–1037.

Fisher, M. L., R. Jaikumar. 1981. A generalized assignment heuristic for vehicle routing. *Networks* **11** 109–124.

Gutek, C., S. Shro□. 2005. Shredding industry has a decent outlook. Tech. rep., Morgan Stanley-Business Services.

Hariharan, R., P. Zipkin. 1995. Customer-order information, leadtimes, and inventories. *Management Science* **41**(10) 1599–1607.

Hemmelmayr, V., K.F. Doerner, R.F. Hartl, M.W.P. Savelsbergh. 2010. Vendor managed inventory for environments with stochastic product usage. *European Journal of Operational Research* **202** 686–695.

Hvattum, L.M., A. Lokketangen, G. Laporte. 2009. Scenario tree-based heuristics for stochastic inventory-routing problems. *INFORMS Journal on Computing* **21**(2) 268–285.

IBISWorld. 2019. *Document Management Services in the US, IBISWorld Industry Report OD5044*. URL `https://clients1.ibisworld.com/reports/us/industry/default.aspx?entid=5044(accessedSeptember3,2019)`.

Iglehart, D. 1963. Dynamic programming and stationary analysis of inventory problems. *Multistage Inventory Models and Techniques, Stanford University Press* **Stanford, CA** 1–31.

Iyer, A.V., J. Ye. 2000. Assessing the value of information sharing in a promotional retail environment. *Management Science* **2**(2) 128–143.

Jaillet, P., J.F. Bard, L. Huang, M. Dror. 2002. Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Science* **36**(3) 292–300.

Ketzenberg, M., M.E. Ferguson. 2008. Managing slow-moving perishables in the grocery industry. *Production and Operations Management* **17**(5) 513–521.

Ketzenberg, M.E., N. Geismar, R. Metters, E. van der Laan. 2013. Value of information for managing retail inventory remotely. *Production and Operations Management* **22**(4) 811–825.

Kleywegt, A.J., V.S. Nori, M.W.P. Savelsbergh. 2002. The stochastic inventory routing problem with direct deliveries. *Transportation Science* **36**(1) 94–118.

Kleywegt, A.J., V.S. Nori, M.W.P. Savelsbergh. 2004. Dynamic programming approximations for a stochastic inventory routing problem. *Transportation Science* **38**(1) 42–70.

Minko, A.S. 1993. A markov decision model and decomposition heuristic for dynamic vehicle dispatching. *Operations Research* **41**(1) 77–90.

Mishra, B.K., S. Raghunathan, X. Yue. 2009. Demand forecast sharing in supply chains. *Production and Operations Management* **18**(2) 152–166.

Prive, J., J. Renaud, F. Boctor, G. Laporte. 2006. Solving a vehicle-routing problem arising in soft-drink distribution. *Journal of the Operational Research Society* **57** 1045–1052.

Ruiz-Benitez, R., M. Ketzenberg, E.A. van der Laan. 2013. Managing consumer returns in high clockspeed industries. *Working paper* .

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

Smartbin. 2009. *Lack of information costs European recycling market 1 billion per year*. URL `http://www.smartbin.com/latest-news-sp-21723/134-lack-of-information-is-costing-european-recycling-market-1-billion-a-year-more-than-necessary.html(accessedApril30,2013)`.

Smartbin. 2010. *Green technology company rebrands to leverage new growth*. URL `http://www.smartbin.com/latest-news-sp-21723/175-green-technology-company-rebrands-to-leverage-new-growth-opportunities.html(accessedApril30,2013)`.

Veinott, A., H. Wagner. 1965. Computing optimal (s, s) policies. *Management Science* **11**(52) 525–552.

Wagner, J. 2008. My 50 years in the shredding industry. *Shredder Industry Report* 2–3.

Yang, W-H., K. Mathur, R.H. Ballou. 2000. Stochastic vehicle routing with restocking. *Transportation Science* **34**(1) 99–112.

Zhang, S.H., K.L. Cheung. 2011. The impact of information sharing and advanced order information on a supply chain with balanced ordering. *Production and Operations Management* **20**(2) 253–267.