European Journal of Operational Research 264 (2018) 1101–1115

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

A Markov decision process model for equitable distribution of supplies under uncertainty

Sefakor Fianu, Lauren B. Davis*

Department of Industrial and Systems Engineering, North Carolina Agricultural and Technical State University, Greensboro, NC 27411, USA

ARTICLE INFO

Article history:
Received 10 July 2015
Accepted 1 July 2017
Available online 9 July 2017

Keywords:
Dynamic programming
Food insecurity
Equity
Markov decision processes
Donations

ABSTRACT

Many individuals suffering from food insecurity obtain assistance from governmental programs and non-profit agencies such as food banks. Much of the food distributed by food banks come from donations which are received from various sources in uncertain quantities at random points in time. This paper presents a model that can assist food banks in distributing these uncertain supplies equitably and measure the performance of their distribution efforts. We formulate this decision problem as a discrete-time, discrete state Markov decision process that considers stochastic supply, deterministic demand and an equity-based objective. We investigate three different allocation rules and describe the optimal policy as a function of available inventory. We also provide county level estimates of unmet need and determine the probability distribution associated with the number of underserved counties. A numerical study is performed to show how the allocation policy and unmet need are impacted by uncertain supply and deterministic, time-varying demand. We also compare different allocation rules in terms of equity and effectiveness.

1. Introduction

Food insecurity is defined as the situation where people are not able to access enough food at all times for an active, healthy life (USDAERS, 2006). The occurrence of food insecurity in the United States (U.S.) is significant, affecting approximately 48.1 million Americans, 15.3 million children and 32.8 million adults. Furthermore, a 2014 study on household food insecurity reported that 14% of U.S. households were unable to access enough food at all times (Coleman-Jensen, Rabbitt, Gregory, & Singh, 2015).

The U.S. Government has established several public assistance programs to address this problem. These programs provide financial assistance, supplemental food, and nutrition support to individuals and states to increase access to healthy food for low-income households. Some of the most well-known programs are the Supplemental Nutrition Assistance Program (SNAP), the Women, Infants, and Children (WIC) program, and The Emergency Food Assistance Program (TEFAP) (see fns.usda.gov/programs-and-services for a comprehensive list). In addition to government funded programs, non-profit organizations, such as Feeding America, also help to address the food insecurity problem.

Feeding America (FA), the nation's leading domestic hunger-relief organization, seeks to eliminate hunger by providing food to those in need through a nationwide network of about 200 member food banks and distribution centers across the country (FeedingAmerica, 2014). FA, as the parent food bank, provides administrative support, training of personnel, standards for food safety and food distribution to its member food banks. The local food banks, although partnering with FA, remain largely independent with their own management systems and budget. They solicit funds, food and supplies from individuals, groups, farmers, local manufacturers and retailers. These donations represent sources of supply that enable them to meet the demand of the people at risk of hunger. Hence, supplies to the food bank are based on the goodwill of donors who are not obligated to give at any particular time or in any particular quantity. Consequently, donations may be infrequent, almost expired which makes them inappropriate for consumption after a few days, or may not be what is actually needed. Nevertheless, food banks need to be able to adequately manage their inventory to ensure equitable distribution of supplies since meeting the demands by aid recipients is not possible with limited supply.

This paper presents a decision model that assists food banks in equitable allocation of uncertain, donated supplies. We define equity as a function of the pounds distributed per person in poverty (PPIP). FA proposed this indicator as a way to measure the performance of its members. The benchmark is to distribute at least 75 pounds of food products for each person in poverty over a 12-month period. We use PPIP to measure equitable distribution practices within the food bank's service area and also identify

* Corresponding author.
E-mail address: lbdavis@ncat.edu (L.B. Davis).

http://dx.doi.org/10.1016/j.ejor.2017.07.017
0377-2217/

deviations from the performance goal set by FA. The decision making problem is formulated as a discrete-time, discrete state (DTDS) Markov decision process (MDP).

The MDP model is used to (i) identify an optimal allocation policy that maximizes equity in the distribution of supplies as estimated by the measured PPIP; (ii) provide estimates of unmet need in each county as a function of the PPIP; and (iii) determine the probability distribution associated with the number of underserved counties. A numerical study is performed using data from the Food Bank of Central & Eastern North Carolina (FBCENC) (2012), a member of Feeding America's network. We investigate three different allocation rules typically used in one warehouse, multiple retailer systems. Our findings show that allocating supply to each county in proportion to its poverty population (proportional allocation) maximizes equity. We also show the relationship between unmet need as a function of changes in supply and demand, characterize the structure of the optimal policy and show how the policy changes based on food donation behavior. While the underlying structure of the optimal decision rule may be somewhat intuitive, our approach to this problem presents some interesting results not explored in the existing literature, particularly as it relates to inventory modeling in humanitarian relief. We modeled the stochastic behavior of the food bank's inventory system using an MDP which has the advantage of indicating the best way to allocate supplies based on the inventory level of the food bank. We present a novel transformation of the state space to account for the large distribution quantities observed in practice and show that the underlying stochastic behavior can be approximated by a normal distribution. We also note that this transformation allows us to easily identify situations where increases in donated supply can allow food banks to further meet their distribution goals. Furthermore, our model considers stochastic supply which is the ideal case for food banks as well as indicates desirable inventory states for the food bank to assist them in proactive planning. In addition, we investigate the current allocation rule the food bank uses and evaluate how far this allocation rule deviates from perfect equity.

The remainder of the paper is outlined as follows. Section 2 summarizes the related literature. Section 3 describes the problem and solution approach. The data analyses and experimental design are presented in Section 4. The results and discussion are summarized in Section 5. Section 6 provides concluding remarks and identifies areas for future research.

2. Related literature

A Humanitarian Supply Chain (HSC) is a network of organizations that ensure the solicitation, transportation, warehousing and distribution of supplies to people affected by emergencies. These emergencies do not only include large-scale catastrophes caused by natural or man-made disasters but also food insecurity primarily caused by economic hardships (Mohan, Gopalakrishnan, & Mizzi, 2013). Due to the increasing trends of natural disasters and food insecurity, HSC management has attracted significant attention (Altay & Green III, 2006; Balcik, Beamon, Krejci, Muramatsu, & Ramirez, 2010; Blanco & Goentzel, 2006; Galindo & Batta, 2013). However, given the importance of inventory management in humanitarian relief operations, the amount of research available in this area is limited compared to that on commercial inventory management.

2.1. Challenges of relief inventories

Relief inventories are referred to as social inventories because they serve broad social objectives as opposed to being used for the benefit of an individual enterprise (Whybark, 2007). Comparing relief and commercial inventory, relief inventories are unique in terms of their source of supplies, objectives, stakeholders, performance measurement and the level of uncertainty and risk (Balcik & Beamon, 2008; Van Wassenhove, 2005). In contrast to commercial inventories, supply in relief inventories is highly uncertain because it is dependent on donations that are constantly evolving. For hunger-relief organizations, the defining sources of supply are scarce government funding and irregular charitable donations from individuals and corporations. Issues with supply uncertainty range from the ability of a donor to give supplies, the varied quantities of supplies donated and the receipt of unsolicited and sometimes unwanted donations (Chomilier, Samii, & Van Wassenhove, 2003).

On the demand side difficulties arise in quantifying the needs for the services of relief organizations. For hunger-relief, food insecurity or poverty levels can serve as an estimate of the demand (Mohan et al., 2013). However, as a result of the limited supplies in hunger-relief operations, unsatisfied demand is very common.

Decision making in inventory management in the presence of random supply and demand can be very challenging with obvious impacts on operating costs and customer service levels. To cope with random supply, commercial inventory managers have adopted sourcing from multiple suppliers (Ahiska, Appaji, King, & Warsing, 2013; Mohebbi, 2003; Tomlin, 2006). Multiple sourcing is also used in relief inventories through the presence of different donors, such as government, corporations, and individuals. However, supply is still frequently insufficient to satisfy the demand. Hence, the objectives of decision making models for relief inventory control usually include finding optimal ways to increase supplies and determining ways to effectively and equitably distribute supplies. Effectiveness refers to the ability to distribute (or allocate) resources in a way that meets the needs of the end customer, whereas equity addresses the ability to allocate resources to multiple customers in a way that is fair (Orgut et al., 2016).

Recent studies in relief inventory management address effective distribution of supplies under supply and demand uncertainty (Rottkemper et al., 2012; Bozorgi-Amiri, Jabalameli, & Mirzapour Al-e-Hashem, 2013; Orgut, Ivy, Uzsoy, & Wilson, 2015). However, studies that emphasize equitable distribution of donated relief supplies are limited. Equitable and effective distribution is paramount in hunger-relief organizations since supply is almost always below demand. Equitable distribution of resources should ensure a fair shairing of the resources among recipients. Only Orgut et al. (2015) and Lien, Iravani, and Smilowitz (2014) consider donated relief inventory distribution, that explicitly incorporates equity. However, neither study directly accounts for uncertainty in supply.

Orgut et al. (2015) develop a linear programming model to determine equitable and effective distribution of donated food for a food bank. Their model maximizes the amount of food distributed while limiting the absolute deviation from a perfectly equitable distribution for each county in the food bank's service area. They also consider the receiving capacities of the counties and develop policies that minimize the amount of undistributed food.

Lien et al. (2014) develop a non-profit sequential allocation model with the goals of equitable distribution of resources and effective service. They define service in terms of fill rate which is calculated as the ratio of the allocated supply to customer demand. The objective function maximizes the expected minimum fill rate among customers, which balances equity in fill rates with effectiveness in the use of resources (low waste).

2.2. Research contribution

This paper also considers equitable distribution of donated supplies. The work of Orgut et al. (2015) is closely related, in the sense that they study equitable distribution of food donations for a food bank. However, we focus on the impact of equity stemming from supply uncertainty rather than agency capacity.

We explicitly incorporate the stochastic nature of donations while modeling the behavior of the inventory system over time. Our framework permits an evaluation of equity and county level predictions of unmet need. Furthermore, we evaluate the improvement in equitable distribution resulting from changes in donation behavior.

Most supply allocation problems in relief inventory are formulated as a multi-objective linear programming problems with objectives such as cost minimization, minimization of travel time, and maximization of satisfied demand (Davis, Samanlioglu, Qu, & Root, 2013; Tzeng, Cheng, & Huang, 2007). Other quantitative decision making models for relief inventory management have been centered on stochastic programming techniques that considered demand uncertainty (Beamon & Kotleba, 2006; Bozorgi-Amiri et al., 2013; Ozbay & Ozguven, 2007). Markov decision processes (MDPs) have not been used to model relief inventories. Most of their applications in inventory management have been centered on commercial inventories to deal with the problem of unreliable suppliers and uncertain demand (Ahiska et al., 2013; Mohebbi, 2003; Tomlin, 2006; Silbermayr & Minner, 2014). This paper uses a DTDS MDP model that considers equitable distribution of supplies as well as the distribution performance of the food bank.

Our work also contributes to the growing literature on supply uncertainty. Supply uncertainty can be classified along three dimensions: timing (i.e. lead time), quantity, and purchase price (Tajbakhsh, Zolfaghari, & Lee, 2007). The research presented in this paper is more closely aligned with uncertain supply quantities. Relevant studies in uncertain supply quantities present strategies that address yield uncertainty or supply interruptions from a production, procurement, or sourcing perspective (Yano & Lee, 1995). The uncertainty in the supply comes from the suppliers being unreliable (Ahiska et al., 2013; Tomlin, 2006; Yang, Aydin, Babich, & Beil, 2012), partially available, unavailable (Güllü, Önol, & Erkip, 1999), or having variable capacity (Tajbakhsh et al., 2007). As a result, the quantity ordered and quantity received are different. However, in the foodbank context, supplies are pushed to the supply chain member based on charitable giving and perceived need rather than in response to a specific order quantity. As such, the uncertainty stems from the behavior of sporadic or one time donors which causes variation in the quantity and quality of the donations. We discuss supply uncertainty within the context of food donations, and explore the effect of supply uncertainty on food distribution decisions. The model can be used for benchmarking the performance of hunger-relief organizations in their efforts to equitably allocate recourses to the people they serve.

3. Problem description and model formulation

3.1. Problem description

The Food Bank of Central and Eastern North Carolina (FBCENC), a member of Feeding Americas' network, has been providing food to people at risk of hunger in 34 counties in Central and Eastern North Carolina for over 30 years. In the fiscal year 2014–15 (FY1415), FBCENC distributed nearly 64 million pounds of food and non-food essentials to aid recipients (FBCENC, 2016). FBCENC operates six branch warehouses in Wilmington, Durham, Raleigh, Sandhills, Greenville, and New Bern. These warehouses collect and sort food donations, conduct quality assessments and then store the supplies prior to distribution. Food donations are distributed to charitable agencies who directly serve individuals in need of food assistance. More than 70% of the supplies to FBCENC are donations from individuals and organizations that are subject to significant uncertainties. Furthermore, the demand that the FBCENC needs to satisfy normally exceeds the supplies that come in.

Each branch receives donations from local sources as well as transfers from other branches. The bulk of the transfers come from the Raleigh branch which serves as the main warehouse and central hub of the distribution network and receives a large amount of donations particularly, those coming from government sources (e.g. TEFAP) and large corporate donors. These donations are subsequently transferred to other branches to ensure fair allocation of high quality donations throughout the network. This is supported by the data from FBCENC which shows that 117,284,024 pounds of donations were received at the Raleigh branch of which 71% were transferred to other branches. In contrast, the Durham branch transferred approximately 8% of food out of their warehouse since they are underserved.

3.2. Model assumptions and sequence of events

Our modeling framework is based on the perspective of a single branch in the food distribution network. We assume that the supply flow into the branch comes from two sources: local donations from the community and transfers from other branches in the network. Rather than considering all types of products, we aggregate all donated product types into a single unit and develop a formulation for a single item inventory system with periodic review. Supply flow out of the branch is based on food need in the service area as determined by the charitable agencies served. In our model, demand from charitable agencies is aggregated to the county level and considered a single demand point. The following assumptions are made about the proposed DTDS-MDP:

1. The state of the system is described by the amount of inventory (in pounds) available at the start of each month;
2. The branch operates a single capacitated warehouse whose capacity determines the upper bound of the state space;
3. Donations are stochastic and occur during the time period;
4. Received donations are added to the current available inventory and are available for distribution during the same time period. This is a reasonable assumption since the time period is long enough for donations to be sorted and reviewed for quality in accordance with the standards for food safety and distribution practices;
5. Food donations are distributed to the counties according to the county demand, available inventory and a predefined allocation rule;
6. Demand in each county is deterministic. This assumption is reasonable since the food bank uses poverty population data from the U.S. Census Bureau to estimate food need in their service area;
7. Demands are filled before the beginning of the next time period;
8. The amount distributed in each county cannot exceed its demand;
9. There is no reallocation of supplies after distribution;
10. The amount transferred from other branches in the network is stochastic and is added to the remaining inventory at the end of the period. A delivery lag is associated with transfers due to the fact that transfers might not reach the receiving branch before the transferring branch has satisfied its demand. The delivery lag can also be attributed to the time associated with processing the transfers before they are available for distribution.

3.3. Model formulation

3.3.1. Decision epoch

The model notation and definitions are summarized in Table 1. Allocation decisions are made on a monthly basis before the beginning of the next month. This is a finite time horizon model and the set of time periods is $T = \{1, 2, \dots, \tau\},\ \tau < \infty$.

Table 1
Model notation and definitions.

Notation		Definition
Sets	V	Set of all possible system states $V = \{\tilde{M}_{LB} \dots \tilde{M}_{UB}\}$
	C	Set of counties to be served $C = \{1, \dots, 6\}$
	A	Set of allocation rules $A = \{a_1, a_2, \dots, a_N\}$
	T	Time periods with $t \in \{1, \dots, \tau\} \tau < \infty$
State variables	v_t	Available inventory at time t (measured in pounds)
Random variables	X_t	Food donations at time t with realization $x_t \in X_t$ (measured in pounds)
	Y_t	Transfers of food from other locations at time t with realization $y_t \in Y_t$ (measured in pounds)
Decision variables	k_{ct}^a	Pounds of food distributed to county c at time t given allocation a
Reward variables	$r_c(v_t, a)$	Pounds of food distributed per person in poverty to county $c \in C$ in state v_t under allocation rule $a \in A$
Other variables	$\tilde{v}_t$	Percentage deviation from mean available inventory at time t
	$\tilde{x}_t$	Percentage deviation from mean donation amount at time t
	$\tilde{y}_t$	Percentage deviation from mean branch transfer at time t
Parameters	P_{ct}	Poverty population in county $c \in C$ at time t
	d_c	Demand for county $c \in C$
	f_c	Fraction of available inventory allocated to county $c \in C$
	H_c	History of total distribution over the previous 11-months to county c
	μ_I	Average inventory in pounds
	μ_D	Average donation in pounds
	μ_B	Average branch transfer in pounds

3.3.2. State of the system

The state of the system is the available inventory which represents the supplies (in pounds) in the warehouse. Based on the data received from FBCENC the available inventory values range from 100,000 to 800,000 pounds. Therefore, the state space is discretized. The discretization procedure is shown in the appendix. The state space is denoted by $V = \{\tilde{M}_{LB} \dots, \tilde{m}, \dots \tilde{M}_{UB}\}$ where $\tilde{m} \in V$ are pseudo states that represent the discretized form of the available inventory values and $\tilde{M}_{LB} < \tilde{M}_{UB}$.

3.3.3. State transitions and transition probability

The events that cause a transition from one state to the next are:

1. Donation, $x_t \in X_t$, which is stochastic with CDF $\Phi_x(\bullet)$
2. Transfers, $y_t \in Y_t$, which are stochastic with CDF $\Phi_y(\bullet)$
3. Distribution of available inventory to aid recipients in county c, k_{ct}^a given allocation decision $a \in A$.

Given the above transition parameters and assuming that donations and transfers are independent events, the available inventory in the next time period v_{t+1} can be computed as a function of the available inventory at time t, v_t, as follows.

$$v_{t+1} = \left[v_t + x_t - \sum_{c \in C} k_{ct}^a \right]^+ + y_t \tag{1}$$

In (1) $v_t = (1 + \frac{\tilde{v}_t}{100})\mu_I$, $x_t = (1 + \frac{\tilde{x}_t}{100})\mu_D$, $y_t = (1 + \frac{\tilde{y}_t}{100})\mu_B$ where v_t, x_t and y_t are the actual values of the available inventory, donation and transfers (in pounds), respectively. The amount distributed will never exceed the available inventory i.e., $\sum_{c \in C} k_{ct}^a \le v_t + x_t$. Therefore, the expression $[v_t + x_t - \sum_{c \in C} k_{ct}^a]^+$will always be non-negative. For ease of notation, let v_{UB} and v_{LB} represent the upper and lower bounds (in pounds) for the available inventory. These values can be easily determined from the pseudo-state variables as $v_{UB} = (1 + \tilde{M}_{UB})\mu_I$ and $v_{LB} = (1 + \tilde{M}_{LB})\mu_I$.

The probability that the system moves from the current state v_t, to the next state v_{t+1}, is influenced by the donation and transfer probabilities and the specific allocation rule a as shown below in equations (2a)-(2d).

$$p(v_{t+1}|v_t, a) \begin{cases} \sum_{x_t} \sum_{y_t \ge v_{UB} - (v_t + x_t - K_t(a))} p(x_t)p(y_t) & v_{t+1} = v_{UB} \quad (2a) \\ \sum_{x_t} \sum_{y_t \le v_{LB} - (v_t + x_t - K_t(a))} p(x_t)p(y_t) & v_{t+1} = v_{LB} \quad (2b) \\ \sum_{x_t} \sum_{y_t = v_{t+1} - (v_t + x_t - K_t(a))} p(x_t)p(y_t) & v_{LB} < v_{t+1} < v_{UB} \quad (2c) \\ 0 & otherwise \quad (2d) \end{cases}$$

It should be noted that $0 \le p(v_{t+1}|v_t, a) \le 1$, $\sum_{v_{t+1}} p(v_{t+1}|v_t, a) = 1$ and $K_t(a) = \sum_{c \in C} k_{ct}^a$. Based on the discretization procedure and given $x_t \in R_m$ or $y_t \in R_m$, $p(x_t)$ *or* $p(y_t) = \phi(R_m^+) - \phi(R_m^-)$, where R_m^+ and R_m^- represent the upper and lower bounds, respectively, of the bin range used in the discretization procedure.

3.3.4. Allocation rules

Allocation rules correspond to the actions that the decision maker (DM) chooses based on the state of the system and, have been widely used in commercial inventory management. An example is the case of 'one warehouse and multiple retailers' (OWMR). Allocation rules identified under the OWMR system similar to what is used in this research are fixed allocation and proportional allocation (Karaesmen, Scheller-Wolf, & Deniz, 2010). For fixed allocation, each retailer receives a predetermined fraction of goods and for proportional allocation, each retailer receives a proportion of goods based on their share of the total demand. For the model considered in this research, the DM uses the allocation rules to distribute supplies to the counties as follows:

Proportional allocation (PA) – Rule 1: Rule 1 uses the proportional allocation approach by Karaesmen et al. (2010) and Orgut et al. (2015). Distribution to counties is such that each county receives supplies based on the ratio of their poverty population to the total poverty population as shown in (3).

$$rule\ a_1: \ k_{ct}^{a_1} = \min\left(\frac{P_{ct}}{\sum_{c \in C} P_{ct}} * (v_t + x_t), \ d_c \right) \tag{3}$$

In practice, FBCENC uses a proportional allocation rule which they refer to as 'fair share' designed to ensure that each county receives food in proportion to its percentage of the overall need. We further consider PA in two unique perspectives: (1) Dynamic Proportional Allocation (DPA) and (2) Fixed Proportional Allocation (FPA). DPA requires updating of the poverty population as the demand changes over time. FPA is a static version of DPA that reflects what may be done in practice. That is, decisions about what to allocate may be based on estimates of demand at a single point in time (e.g. census estimates) even though demand changes over the time horizon.

Serve largest demand first (SLDF) – Rule 2: With SLDF, the DM serves the county with the largest demand first as shown in (4). Based on what is left after the previous distribution, the DM proceeds to serve the next larger demand and eventually serves the least demand last. However, under this approach inventory may run out before serving the county with the smallest demand.

$$\text{rule } a_2: \; k_{ct}^{a_2} = \min\left(\max\left(\upsilon_t + x_t - \sum_{i \in F_t^*} k_{it}, 0\right), d_c\right) \tag{4}$$

where, $F_t^* = \{c' \in C | d_{c'} > d_c\}$. F_t^* is a set of all previous deterministic demands $d_{c'}$ that have been served.

Serve smallest demand first (SSDF) – Rule 3: For SSDF, the DM serves the county with the smallest demand first as shown in (5). Based on what is left after the previous distribution, the DM proceeds to serve the next smaller demand and eventually serves the highest demand last.

$$\text{rule } a_3: \; k_{ct}^{a_3} = \min\left(\max\left(\upsilon_t + x_t - \sum_{i \in F_t^*} k_{it}, 0\right), d_c\right) \tag{5}$$

In (5), $F_t^* = \{c' \in C | d_{c'} < d_c\}$. F_t^* is a set of all previous deterministic demands $d_{c'}$ that have been served.

3.3.5. Reward determination

The goal the food bank is to equitably distribute supplies to aid recipients through its warehouses and achieve a long-term goal of meeting the PPIP target of 75 over a 12-month period. Equity in distribution is defined to ensure that each person in poverty in the service area receives an equal share of the food distributed. The reward is therefore an objective function that maximizes equity in distribution as measured by the PPIP.

Pounds distributed per person in poverty (PPIP): The PPIP for a county is the ratio of the current distribution plus the previous 11 months distributions to the poverty population in that county as shown in (6). The previous 11 months distributions is an estimate of distributions for previous 11 months that is initialized once at the beginning of the process for each county and used throughout the time-horizon.

$$r_c(\upsilon_t, a) = \frac{k_{ct}^a + H_c}{P_{ct}} \tag{6}$$

We define $PPIP_t$ as the target PPIP as set in accordance with the Feeding America performance indicator as the benchmark to measure the performance of the food bank branches. Thus, if $r_c(\upsilon_t, a) < PPIP_t$ the county is considered to be underserved in period t. If $r_c(\upsilon_t, a) > PPIP_t$, the county is said to be over-served in period t. Otherwise, the county is well served.

Measure of equity: There are several metrics that are used to measure equity, including the difference between the maximum and minimum values, variance, coefficient of variation, sum of absolute deviations, maximum deviation, and mean absolute deviation (Marsh & Schilling, 1994). Equity is maximized by minimizing these measurements. We formulate a measure of equity based on the mean absolute deviation of the PPIP (Δ_{PPIP}) from the mean PPIP for all counties as shown in (7) and (8).

$$\bar{r}(\upsilon_t, a) = \frac{1}{|C|} \sum_{c \in C} r_c(\upsilon_t, a) \tag{7}$$

$$\Delta_{PPIP}(\upsilon_t, a) = \sum_{c \in C} \frac{|r_c(\upsilon_t, a) - \bar{r}(\upsilon_t, a)|}{\bar{r}(\upsilon_t, a)} \tag{8}$$

The greater the value of Δ_{PPIP} the less the equity; perfect equity is achieved when $\Delta_{PPIP} = 0$. Thus, minimizing (8), maximizes equity.

Expected immediate reward: The immediate reward, which is the expected reward for state υ_t under allocation rule a is shown in (9), where the expectation is taken with respect to the donations X_t and the transfers Y_t.

$$\boldsymbol{q}(\boldsymbol{\upsilon_t}, \boldsymbol{a}) = \boldsymbol{E}_{\boldsymbol{X_t}, \boldsymbol{Y_t}}[\Delta_{\boldsymbol{PPIP}}(\boldsymbol{\upsilon_t}, \boldsymbol{a})] \tag{9}$$

3.3.6. Additional measures of distribution effectiveness

In addition to measuring equity, we also compute the expected unsatisfied demand, and the probability of counties being underserved. Unsatisfied demand is the amount (in pounds) of additional supplies needed by the counties to meet the target PPIP ($PPIP_t$) and can be estimated for each inventory state as $E_{X_t, Y_t}(PPIP_t - r_c(\upsilon_t, a))$. The probability distribution of the number of underserved counties is determined according to the law of total probability illustrated in (10), where p_s denotes the steady state probability, $p(N_c = n)$ is determined with respect to the uncertain supply and $I(\cdot)$ is an indicator function that takes the value of one if $PPIP_t - r_c(\upsilon, a)$ is positive and zero otherwise. The dependence on x is implied through the allocation amount k_{ct}^a defined in (3)–(5).

$$p(N_{underserved} = n) = \sum_{\upsilon \in V} p(N_c = n | \upsilon) p_s(\upsilon) \tag{10}$$

where $p(N_c = n) = \sum_x I(PPIP_t - r_c(\upsilon, a)) p(x)$

3.4. Model evaluation

We seek an optimal policy that maximizes equity. We are interested in both the short term and long term behavior of the system. Therefore, two solution approaches are used. We first use the backward induction approach (Puterman, 2009) to solve the finite horizon (12-month) problem. The policy-iteration approach (Howard, 1960) is used to determine the long term behavior of the model over an infinite time horizon. We are interested in the steady state probabilities of the inventory states and average inventory level which are typical measures of performance for commercial inventory management problems.

4. Data analysis and experimental design

Data was obtained from FBCENC transactions records from 2006/2007 to 2013/2014 representing eight fiscal years. The data is aggregated on a monthly basis resulting in 96 observations representing the pounds of food received for each month. The Durham branch of the food bank is investigated with a focus on dry goods (all items classified under dry storage type).

4.1. Demand estimation

The Durham branch serves Chatham, Durham, Granville, Orange, Person and Vance Counties in North Carolina. The poverty populations of these counties are obtained from the FBCENC's fair share program (FBCENC, 2012). Durham County has the largest poverty population, which is approximately 44% of the entire poverty population being served by the Durham branch.

FBCENC does not have records of the demand for each county. However, we assume that there exists a correlation between the demand of a county and the poverty population of that county (Wight, Kaushal, Waldfogel, & Garfinkel, 2014). Based on the poverty population and the PPIP criterion set by FA (PPIP should be 75 over a 12-month period), the demand by county can be estimated as shown in (11). Table 2 shows the estimated poverty population of each county served by the Durham branch and their projected monthly demands.

$$\text{Projected monthly demand} = \frac{(P_c \times 75)}{12} \tag{11}$$

Table 2
Counties poverty populations and monthly projected demands.

County	Poverty population	Projected monthly demand (pounds)
Chatham	8028	50,175
Durham	36,504	228,150
Granville	5770	36,063
Orange	16,475	102,969
Person	5829	36,431
Vance	10,859	67,869
Total	83,465	521,657

Table 3
Summary of the percentage deviations of each dataset.

Variable	Minimum percentage deviation (%)	Maximum percentage deviation (%)
Available inventory	-53	114
Donations	-77	159
Transfer in	-89	124

4.2. Supply data transformation

The values of the available inventory, donations and branch transfer data are continuous variables. Consequently, these values are discretized using the procedure described in the appendix in order to use them in the DTDS MDP model. The discretization transforms each data point in the original series from pounds into deviations from the mean. More specifically, given a data point s_t, the transformed data point $\tilde{s}_t$is equal to $(s_t - \bar{s})/\bar{s}$, where $\bar{s}$ is the mean for the series. Figs. A.1–A.3 in the appendix show the transformed series for available inventory, donations, and branch transfers, respectively. The maximum and minimum values are shown in Table 3.

4.3. Probability distributions

The donations and branch transfers are two stochastic events that cause the available inventory to transition from one state to another state. The JMP statistical software was used to fit probability distributions to the donation and transfer mean percentage deviations. The Gaussian probability distributions provided the best fit for both datasets with their respective mean and standard deviations as follows; the transformed donation dataset is normally distributed with mean -4.51 and standard deviation 35.30, and transformed transfer dataset is normally distributed with mean -2.56 and standard deviation 31.98. It should be noted that the mean of the normal distributions for donation and the transfer data sets (-4.51 and -2.56, respectively) are percentage deviations. A goodness of fit test was conducted using the Shapiro–Wilk W Test at 5% significance level. Subsequently their p-values were greater than the significance value. This confirmed that at the 5% significance level, there was enough evidence to conclude that both the donation and transfer datasets were from a normal distribution.

Additional statistical tests were conducted to evaluate: the stationarity of the donation and transfer transformed data using the Augmented Dickey–Fuller (ADF) test (Appendix A3) and stationarity of the transition probabilities using the Anderson and Goodman (1957) approach (Appendix A2). At a 0.05 significance level, the donation and transfer transformed data were stationary since their absolute test statistics were greater than the absolute critical value of 2.93 (Sjö, 2008). The hypothesis test (Appendix A2) verified our assumption that the transition probabilities are stationary over the time horizon.

Table 4
Summary of input parameters.

Parameter name	Notation	Value
State space	V	$\tilde{M}_{LB} = -50\%,\ \tilde{M}_{UB} = 90\%,$
Number of counties	C	6
Target PPIP	$PPIP_t$	75
Set of allocation decisions	A	{1,2,3}
Finite-time horizon (months)	τ	12
Average inventory (pounds)	μ_I	418,000
Average donation (pounds)	μ_D	129,000
Average transfer (pounds)	μ_B	289,000

4.4. Previous 11-month distributions

The prior 11-month distribution records are very important in this research because they are used in the PPIP calculation in (7). In practice, these values would be updated on a rolling horizon and therefore change as a function of monthly distribution policies. In our model, we assume that the 11-month distribution history is stationary. Future work will explore the impact of history on metrics of distribution performance.

4.5. Experimental design

A computational study is performed to analyze the optimal allocation policy, the unsatisfied demand and the number of underserved counties for the Durham branch. The model is evaluated using different input parameters in order to answer the following research questions: (1) Should a fixed allocation policy be used at all times? (2) Can an allocation policy be defined generically for different demand cases? (3) How does a large influx of supplies influence the allocation policy or how does very low supplies influence the allocation policy?

4.5.1. Base scenario

A base scenario is established to gain insight into the optimal policy structure using the projected monthly county demands. We define this as demands needed by the counties in order to meet the objective of distributing 75 pounds of food per person over a 12-month period. The results from the base scenario form the baseline from which all other scenarios are compared.

4.5.2. Sensitivity analysis

Changing supply and demand: The demands and the supply values are varied to generate different cases to test the behavior of the model outputs. All results are presented as deviations from the base scenario.

$$Error(deviation) = Measured\ result - Base\ result \tag{12}$$

In order to understand the impact of using a static allocation policy at all times, two demand scenarios are investigated: (i) stationary county demands over the entire 12-month period, as defined in Table 2 and (ii) non-stationary demand which represents a case where county demands remain unchanged for the first six months and then fluctuate for the remaining six months. For each of these two demand scenarios, different cases are generated by adjusting each poverty population and projected monthly demand from -50% to 100% in increments of 10%.

To analyze the effect of changes in the supply, four parameters are changed, one at a time, from -50% to 50% in increments of 10%. These parameters are donation mean, donation standard deviation, transfer-in mean and transfer-in standard deviation. This type of experiment may highlight the impact of increasing efforts to solicit more donations due to unsatisfied demand. Table 4 summarizes the input parameters used in the model.

Table 5
Previous 11-month distributions to each county.

County	Previous 11-month distributions (pounds) for scenario 1	Previous 11-month distributions (pounds) for scenario 2
Chatham	551,925	238,775
Durham	2509,650	1334,602
Granville	396,688	251,217
Orange	1132,656	435,702
Person	400,744	175,522
Vance	746,556	366,906

Effect of distribution history (previous 11-month distribution): To analyze the impact of the distribution history on the model output, two scenarios are investigated: 1) county demands are satisfied during all the previous 11-month. In other words, the previous 11-month distribution meets the demand for a given county; 2) county demands are not satisfied all the time during the previous 11-month. The first scenario in Table 5 represents the previous 11-month distribution, which was obtained by multiplying the projected monthly demand by eleven for each county. For the second scenario, the previous 11-month distribution is obtained from the historical data provided by the food bank.

5. Results and discussion

5.1. Base scenario

5.1.1. Finite horizon analysis

Table 6 describes the results for the finite-horizon problem. The average food distribution performance metrics and optimal allocation rule are defined according to the pseudo-state and actual state values (in pounds) The optimal policy that achieves perfect equity is stationary and has the following form.

$$\pi^*(\tilde{v}_t) = \begin{cases} 1 & \tilde{v}_t \le 8\ (1.15\mu_i)\ \forall t \\ \{1,2,3\} & \tilde{v}_t > 8\ \forall t \end{cases} \tag{13}$$

This policy is a threshold policy that partitions the state space into supply constrained and supply abundant states. In supply constrained states, equity can only be achieved when supply is allocated in proportion to the county demand (rule 1-PA) this result is similar to the work of Orgut et al. (2015). We show the optimality of rule PA in Appendix A.4.2. Any state value that is 15% or more above the historical sample mean inventory level (μ_I) is an abundant state, and therefore any demand-based allocation rule (PA, SLDF, SSDF) is optimal. Since the total demand for all six counties is 521,700 pounds, there is clearly enough supply in states 9 through 16 to satisfy all the county demands.

Although it is intuitive that perfect equity can be achieved, an equitable solution does not imply that all demand is satisfied. On average, all counties are underserved unless the inventory level exceeds 15% above the mean. There are a few donation and transfer scenarios in the supply constrained states (states 5 through 7) where some counties are being well served, thus leading to a positive expected number of underserved counties (Table 6). However, all counties are consistently well served in supply abundant states.

A further examination of distribution performance at the county level shows that the food bank can expect unsatisfied demand at the end of each time period to range from 13 to 21 pounds per person in poverty over a 12-month time frame (Table 6). This measure of distribution performance serves as an indicator of the amount below the target PPIP that is realizable. We must note that our results reflect only one product type and are not indicative of the entire capability of the food bank network. However, we illustrate that our modeling framework can identify whether the food bank is meeting or exceeding the FA target.

We also note that it is possible to have fewer counties underserved or a smaller total unmet need per person in poverty when using the other allocation rules (SSDF, SLDF). It is intuitive that SLDF approach prioritizes distribution to larger counties, whereas SSDF prioritizes distribution to smaller counties. Thus, some counties will have all demand met at the expense of the other counties. In addition, SSDF can lead to smaller overall unmet need (see Appendix A.4.3 for a proof regarding this result). However, this behavior is not appealing in the context of food aid relief since neither SSDF nor SLDF ensure a fair allocation of available donations. The proportional allocation approach ensures unmet need is the same across all counties.

5.1.2. Infinite horizon analysis

Since the optimal policy for the finite horizon is stationary, solving the average reward infinite horizon case is not necessary. We can simply compute steady state performance metrics from the Markov chain induced by the optimal allocation policy. Nevertheless, we apply the policy iteration method with our transformed reward structure to confirm results of the finite horizon model. As before, the optimal policy that achieves equity is the proportional allocation rule. A number of interesting results emerge from the steady state probabilities listed in Table 7. Our model suggests that the food bank branch is operating in a supply constrained en-

Table 6
Results for base scenario.

Pseudo states		Actual state (pounds)	Optimal allocation rule(s)	Average food distribution performance metrics over 12-month time horizon		
				Equity*	Expected unmet need (PPIP) per county	Number of counties underserved
<-50%	1	209,000	1	0	21.81	6
-45%	2	229,900	1	0	21.81	6
-35%	3	271,700	1	0	21.05	6
-25%	4	313,500	1	0	20.60	6
-15%	5	355,300	1	0	20.00	5
-5%	6	397,100	1	0	19.43	3
5%	7	438,900	1	0	18.68	1
15%	8	480,700	1	0	18.07	0
25%	9	522,500	1,2,3	0	17.46	0
35%	10	564,300	1,2,3	0	16.87	0
45%	11	606,100	1,2,3	0	16.25	0
55%	12	647,900	1,2,3	0	15.64	0
65%	13	689,700	1,2,3	0	15.03	0
75%	14	731,500	1,2,3	0	14.44	0
85%	15	773,300	1,2,3	0	13.88	0
>90%	16	794,200	1,2,3	0	13.59	0

Table 7
Summary of base scenario results for the infinite horizon.

States	Steady state probabilities
1	0.1796
2	0.193
3	0.1251
4	0.2159
5	0.0978
6	0.0738
7	0.0721
8	0.0216
9	0.0141
10	0.0047
11	0.0014
12	0.0006
13	0.0002
14	0.0001
15	0
16	0

vironment 97% of the time with an average monthly inventory of 301175.2 pounds (corresponding to an expected pseudo state value of approximately 28% below the mean). The expected number of underserved counties is 5 with probability distribution shown in (14). Due to the fact that the optimal policy results in perfect equity, there are only two possible realizations for underserved counties, 0 and 6. Our results help to confirm what is true in practice: supply is significantly less than demand and thus it is important to allocate donated food in an equitable manner.

$$P(N_{underserved} = n) = \begin{cases} 0.16 & n = 0 \\ 0.84 & n = 6 \\ 0 & otherwise \end{cases} \tag{14}$$

5.2. Sensitivity analyses

As with any inventory model, the behavior with respect to supply and demand is quite intuitive. Increasing available supply results in lower unsatisfied demand. However, it is interesting to explore the change in performance relative to percentage increases and decreases in supply and demand. This provides a way to communicate effectively with decision makers about the impact of distribution decisions. The following sections summarize the changes in the optimal policy, unsatisfied demand and the number of counties underserved as a function of changes in the donation behavior (supply) and food need (demand). Results associated with the finite-horizon and infinite-horizon are presented.

5.2.1. Effects of supply uncertainty

Structure of optimal policy: Because of the nature of our problem (stochastic donations and deterministic demand), we are able to develop a closed form expression for the relationship between the policy structure and donation behavior. Eq. (15) represents the percentage adjustment (α_v) in the average sample mean donation that is required in order to satisfy the total demand in state v, thus changing the state from constrained to abundant. In (15), the numerator (when positive) represents the unmet need, and the denominator the smallest realizable case in terms of incoming donation quantities. (The derivation of this quantity can be found in Appendix 4.1).

$$\alpha_v = \frac{(\sum_{c\in C} d_c - v)}{\mu_D * (1 + \tilde{x}_{LB})} - 1 \tag{15}$$

Fig. 1 shows the relationship between the unmet need percentage and the number of constrained states. For example, if the donation sample mean increased by 5.83%, the number of supply constrained states would decrease by one. Extreme changes in the donation quantity are needed to change the structure of the optimal policy (i.e. shifting constrained states to supply abundant states within the current bounds of the state space). Eq. (16) also shows that the structure of the optimal policy is insensitive to changes in the donation standard deviation as well as the transfer-in standard deviation. However, this is not the case when we examine the effect of supply uncertainty on the distribution performance measures. This behavior is discussed in the following sections.

Unsatisfied demand: While changing the supply has a small effect on the structure of the optimal policy, it does affect the level of unsatisfied demand (or unmet need) more significantly. These results are more evident when unsatisfied demand is examined relative to the sample coefficient of variation as shown on Fig. 2. The results are generated by varying the parameters (mean and standard deviation) associated with the normal distribution for the transformed donation and transfer-in data sets as described in Section 4. The values for the coefficient of variation are negative because the mean for both the transformed donation and transfer-in data sets is negative (refer to Section 4.3). The results are shown for states 1, 8 and 16 which represent the lowest, average, and highest inventory level of the food bank, respectively. The results at the initial time period are presented as this encompasses the expected reward over the 12-month time horizon given the starting inventory state. The deviation in the unsatisfied demand for each state increases monotonically as the coefficient of variation increases. Positive deviations indicate higher unmet need whereas

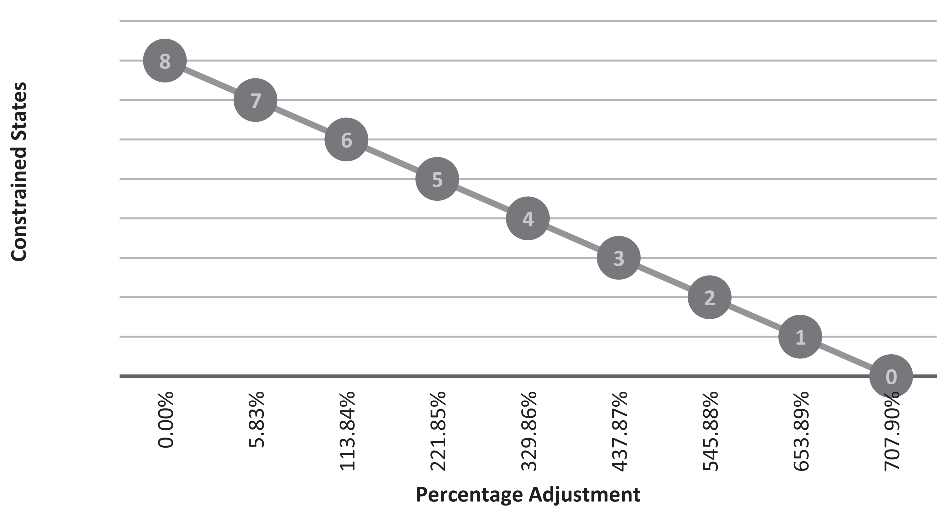

Fig. 1. Optimal policy structure as a function of α_v.

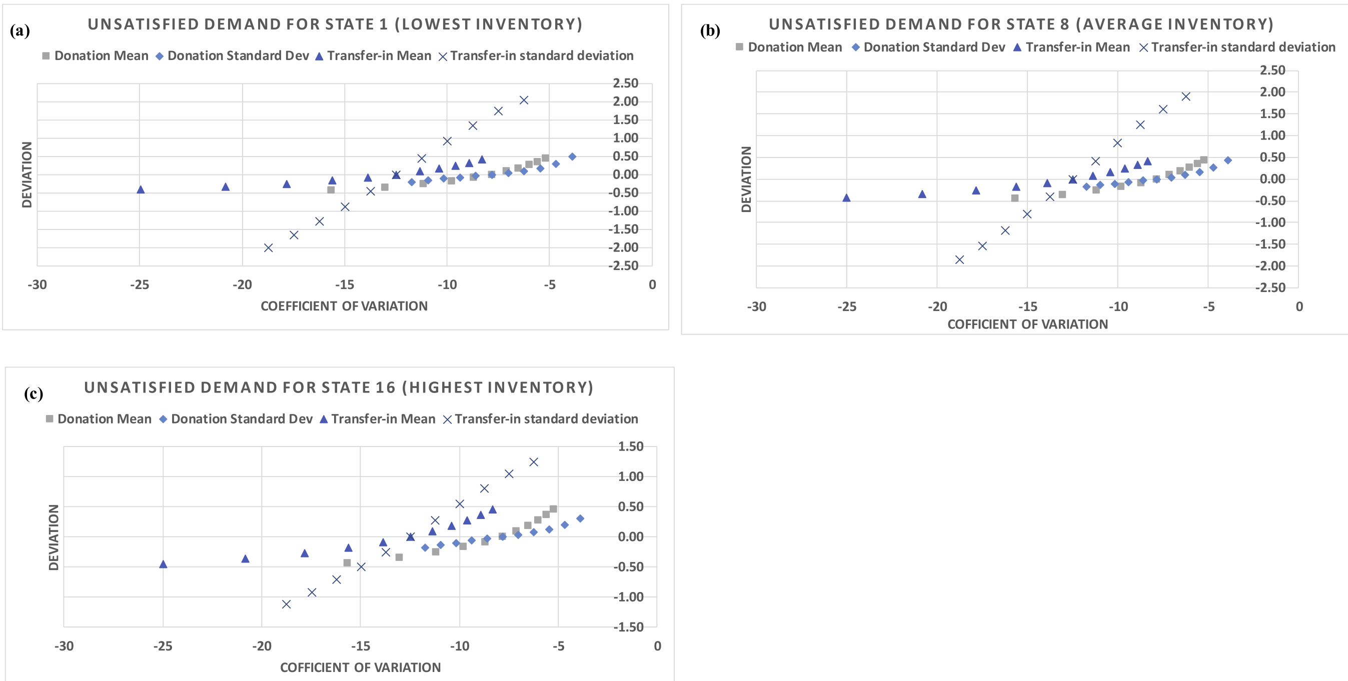

Fig. 2. Deviation of unsatisfied demand relative to coefficient of variation (finite horizon results).

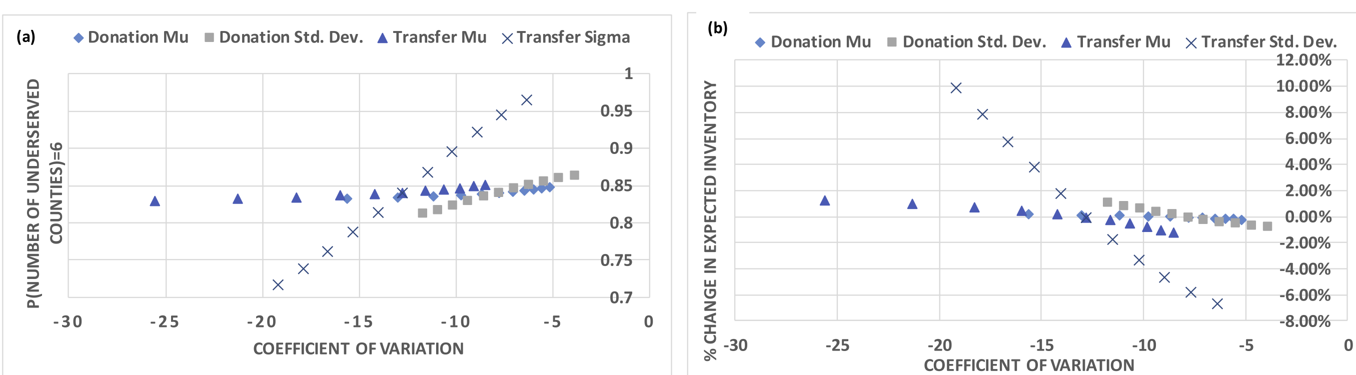

Fig. 3. Long-run (a) probability that all counties are underserved and (b) percentage change in average inventory level from base case.

negative deviations indicate the opposite. More specifically, as supply uncertainty increases, unsatisfied demand increases which increases the probability that all counties are underserved (Fig. 3(a)). This behavior is a result of the overall decrease in the long-run average inventory level (Fig. 3(b)). Furthermore, changing the standard deviation associated with transfers has the most pronounced effect on the ability of the food bank to satisfy unmet need.

We also explored the effect on the ability to satisfy demand by changing the donation and transfer-in sample mean (while keeping the distribution parameters fixed) as shown in Fig. 4. The results show that changing the donation sample mean can increase (decrease) the unmet demand by as much as 9 pounds per person in poverty if the supply is decreased (increased) by 50%. The effect on unmet demand is even higher for changes to the mean amount transferred to the warehouse from other branches. This is because the mean transfer-in is approximately 289,000 pounds, nearly twice as much as the donations.

For this particular data set, the results of our analysis indicate that all six counties can still be well-served under two conditions: (i) if the average donation amount decreases by no more than 20%; and (ii) uncertainty about the true donation behavior, given the current average mean donation amount is unchanged, increases by no more than 30%. Percentage adjustments outside of these two bounds increase the number of underserved counties.

We conclude our analysis of supply uncertainty by stating some observations with respect to traditional inventory models. In traditional inventory models with supply disruption, increasing supply uncertainty results in higher order/production quantities (Güllü et al., 1999; Tomlin, 2006). Another impact of increasing variability in the supply is sourcing from multiple suppliers where less quantities are ordered from the unreliable supplier with a lower price (Ahiska et al., 2013). This results in an increase in the overall inventory cost (Begen, Pun, & Yan, 2016; Mohebbi, 2004; Yeo & Yuan, 2011). However, in our paper, higher uncertainty in supply results in a decrease in average inventory since most of the supply is donated. This results in higher unsatisfied demand as well as the number of underserved counties.

5.2.2. *Effects of demand uncertainty*

Stationary demand: One of the key assumptions of this paper is that of deterministic demand. While this assumption is valid given the context of our study, it is possible for estimates of demand to be imprecise and inaccurate. As shown in Fig. 5, if demand increases by more than 50%, all inventory states are supply

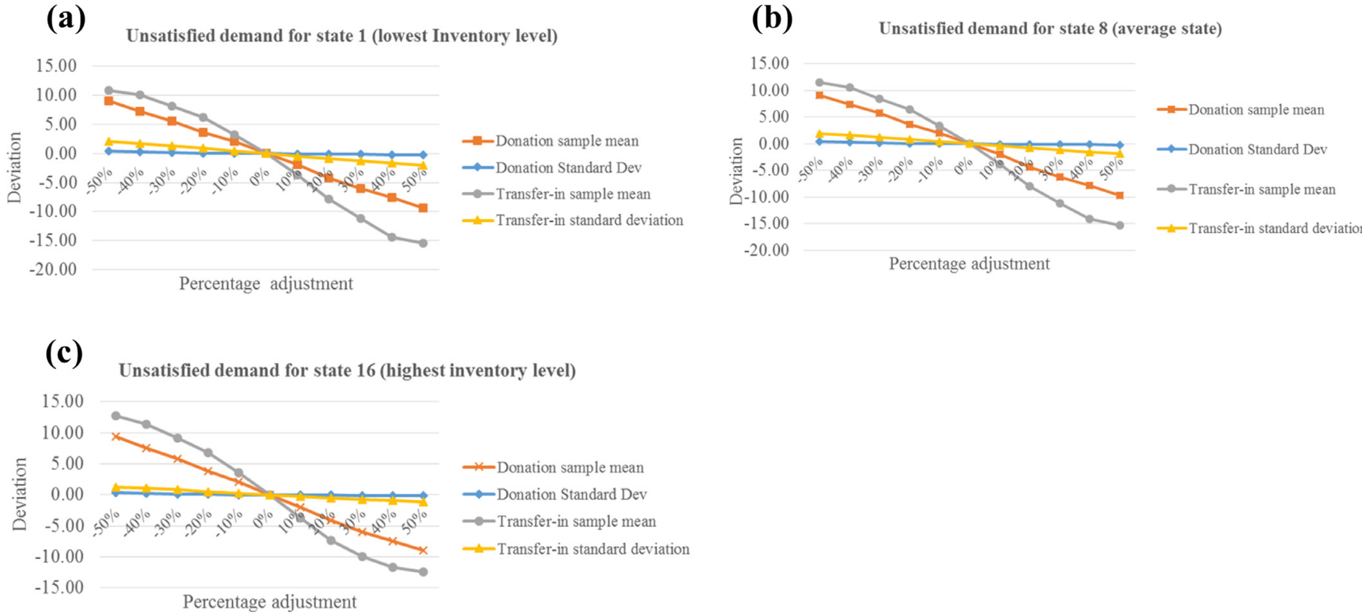

Fig. 4. Deviation of unsatisfied demand for donation and transfer-in sample mean.

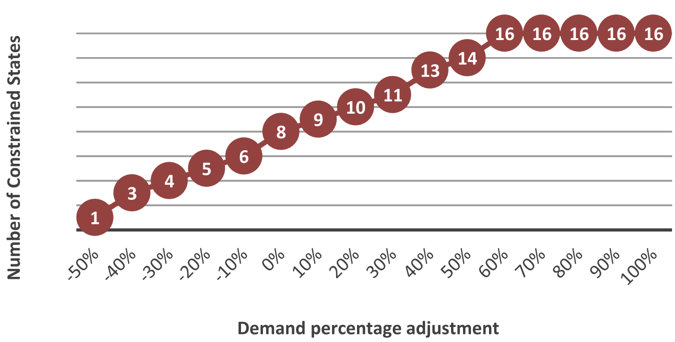

Fig. 5. Optimal policy for stationary demand cases.

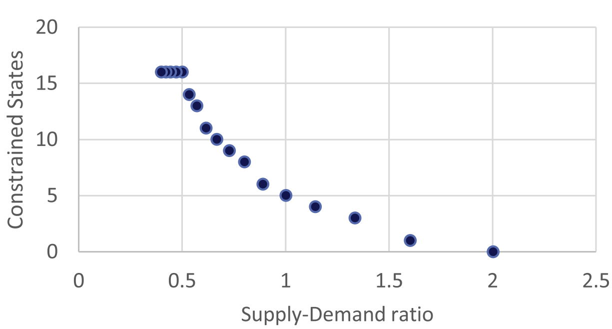

Fig. 6. The number of constrained states as a function of the supply-demand ratio.

constrained, indicating the need for distributing the uncertain donation supplies equitably. It is also interesting to note that even a 10% deviation in the demand estimate can impact the number of supply constrained states, both favorably and unfavorably. Fig. 6 sheds more light on this behavior within the context of the ratio of supply to demand (R_{sd}), measured in pounds. Supply is determined as the sum of the sample mean donation and transfer in amount. Demand in pounds is determined relative to the base amount and percentage adjustment (α_D). Essentially, when this ratio reaches the value of 2, there are no constrained states.

$$R_{sd} = \frac{\mu_D + \mu_B}{\sum_{c\in C} d_c(1+\alpha_D)} \tag{16}$$

Fig. 7 provides a perspective of unmet need in each county. As expected, as demand increases the supply becomes insufficient to satisfy most of food need in each county, resulting in higher deviations from the base case. However, unsatisfied demand does not increase linearly with percentage adjustments in food need. Small deviations in the demand estimate (±10%), results in unsatisfied demand below 7 pounds per person in poverty. In addition, small deviations yield the same results as the base case in terms of the number of underserved counties. However, percentage increases in demand above 10%, increase the number of underserved counties for the average inventory case (state 8). In particular, all the counties are underserved if the demand increases above 30%.

Non-stationary demand: In addition to exploring changes in the total demand, the assumption of stationary demand is relaxed. The intent is to explore the impact of using a fixed proportional allocation rule, given demand changes over time. It is possible for example, for food banks to use the DPA rule assuming demand is not changing over time. When the demand changes, and the allocation quantities are not changed accordingly, the food bank is in effect using a fixed proportional allocation (FPA) rule. Fig. 8 shows the deviation in unmet need between the DPA and FPA rule for state 1 (supply constrained). Even a 10% adjustment in demand can cause an expected deviation of almost 500 pounds per person in poverty considering all counties under the FPA.

5.4. Sensitivity analysis for the distribution history

All the results above assume that county demands are satisfied during the previous 11-months. In this scenario, the optimal policy is rule 1 for the supply constrained states and policies 1, 2, or 3 are optimal for the supply abundant states. However, for the case where all county demands are not satisfied during the previous 11-months the results show higher unsatisfied demand and all counties are underserved irrespective of the state.

6. Conclusion

In this paper, we develop a discrete-time, discrete space Markov decision process model to assist food banks in distributing supplies equitably as well as measure their performance using the PPIP indicator proposed by FA. The modeling approach presented

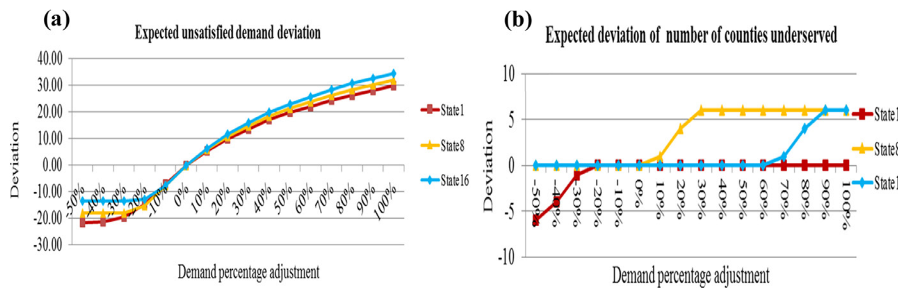

Fig. 7. Deviation of unsatisfied demand and the number of counties underserved.

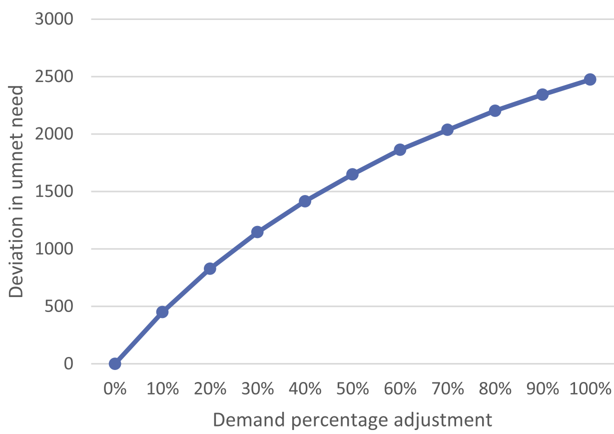

Fig. 8. Deviation of unsatisfied demand between DPA and FPA.

in this paper has some interesting results not explored in the existing literature, specifically as it relates to incorporating supply uncertainty in relief inventory models. We modeled the food bank's inventory system using MDPs which has the advantage of indicating the best way to allocate supplies based on the inventory levels (states) of the food bank. We found a unique approach to represent the large, continuous inventory levels of the food bank that can help identify desirable inventory states to assist them in proactive planning. Our model investigates the impact of donations and transfers on the constrained states and identifies the amount of additional donations needed to move to a completely unconstrained inventory state. Our model also provides some bounds on the deviations from equity that can occur when the optimal allocation rule is not selected. We further describe bounds on unmet need as a function of changes in the average supply and the variation associated with the supply estimates.

From this research, we found that the optimal supply allocation policy that maximizes equity in the distribution of supplies to counties using the PPIP criterion in general is as follows:

1. The proportional allocation rule should be used if the available inventory falls by at most 15% below the average available inventory irrespective of the time period;
2. Any allocation rule can be used if the available inventory is at least 25% above the sample mean irrespective of the time period;
3. Exception: allocation rule 1 should be used throughout the time period if the total county demand exceeds the available inventory to ensure equity.

6.1. Implication for food bank operations

Based on the proposed model and the experimental analysis, our model provides guidance on how to set inventory targets (and indirectly donation targets) in order to ensure counties are well served. In the absence of being able to meet those target inventory levels, the policy obtained from the MDP model indicates how food distribution efforts should be done in order to ensure each county gets their fair share. We show that it is also helpful to update the fair share estimates as demand information changes. While obtaining the information on true demand may be challenging, our results indicate that there could be some significant deviations in unmet need if a fixed proportional allocation rule is used rather than a rule that adapts to changing demand (e.g. DPA). When adopted, the model's input parameters; supply sample means and standard deviations as well as the demand should be updated from time to time so that optimal distribution policies can be updated as information changes.

Our results also describe the relationship between uncertain supply and average inventory levels. More specifically, as supply uncertainty increases, average inventory levels decrease. In traditional inventory models, this behavior can be mitigated with diversification of supply sources or increasing production/procurement quantities. However, these strategies are not applicable for food banks since the supply is mostly donated. This suggests that other approaches to influence donations or supplement potential inventory shortages must be investigated.

6.2. Future work

Irrespective of the work that has been done in this paper, there is still room for improvement. The model could be extended to all other branches in the FBCENC network to study branch to branch variability. One may model the donation and transfer-in with other probability distributions and analyze the associated prediction errors. Likewise, another parameter such as transfer-out can also be added and modeled, perishable items can be considered since this research only investigated the case of dry goods. Warehouse capacity constraints can also be investigated to see how that may affect the optimal policy. Furthermore, a continuous state Markov decision process can be investigated to avoid discretization and the errors that may be associated with it.

We also know that in practice, while the proportional allocation rule may be optimal to achieve perfect equity, there are other constraints not incorporated in the model that may make this difficult to achieve. For example, as mentioned in Orgut et al. (2015), capacity constraints at the agency level may cause deviations from equity which cannot be avoided. In our work, we have not considered capacity constraints at the agencies. The incorporation of capacity and supply constraints is an interesting area of future research.

Acknowledgments

We would like to express our sincere appreciation to Charlie Hale and Earline Middleton from the Food Bank of Central and Eastern North Carolina for their valuable input into this project. This research was supported by NSF grants CMMI-1000018 and CMMI-1000828. We would also like to express our appreciation to the anonymous referees, whose comments helped improve the exposition of this paper.

Appendix

A.1. Discretization procedure

1. Mean percentage deviation

A heuristic approach called the mean percentage deviation (MPD) shown in (A.1.1) below is used to calculate the percentage deviation of the actual value α from the mean value, μ.

$$MPD = \frac{\alpha - \mu}{\mu} \times 100\% \tag{A.1.1}$$

Thus, the donation, branch transfer and available inventory continuous data can be represented with their actual means $\mu_D,\ \mu_B,\ \mu_I$ respectively and their mean percentage deviations.

2. Binning

The binning technique is used to group the MPD values into bins with equal-width. Let α_{min} and α_{max} be the minimum and the maximum percentage deviation of the actual values. Thus the set of percentage deviation of the actual values α is bounded by the range $\alpha_{min} \le \alpha \le \alpha_{max}$ where $\alpha_{min} > -\infty$ and $\alpha_{max} < \infty$. The percentage deviation of the actual values α are grouped into bins taking into consideration α_{min} and α_{max}. Let M be the number of bins, which are numbered, 1 through M, $m \in M$. Also let Δ_α be the bin width given by $\Delta_\alpha = \frac{(\alpha_{\max} - \alpha_{min})}{M}$. Then, the range, R_m, of the mth bin is as shown in (A.1.2).

$$R_m = (\alpha_{min}(m-1)\Delta_\alpha \ \ \alpha_{min} + m\Delta_\alpha] \tag{A.1.2}$$

It should be noted that, the choice of the number of bins M is discretional. The lower and the upper ranges for R_m are given by $R_1 = (-\infty \ \ \alpha_{min} + \Delta_\alpha]$ and $R_M = [\alpha_{min} + (M-1)\Delta_\alpha \ \ \infty)$, respectively. This boundary ranges are essential to cater for unknown data points that might fall outside the predefined domain, $[\alpha_{min} \ \ \alpha_{max}]$ during the lifetime of the model. In our approach, the values to bin are the percentage deviations from the mean and the bin width is 10% for each dataset. Consequently, each percentage deviation value belongs to one of the bins.

The percentage deviations of the available inventory, donation and branch transfer are grouped into bins of equal width. The bins are associated with distinct discrete values using a one-to-one mapping. The median value of each bin range is used to calculate the actual values of the available inventory, donation and transfer with exception of the extreme values. Table 4 shows the upper and lower bounds of the discretized percentage deviations and the number of bins of each dataset. The form of the lower and upper bound is <-50 or $90>$, respectively. The dimensions of the state space are defined based on the bounds for the available inventory.

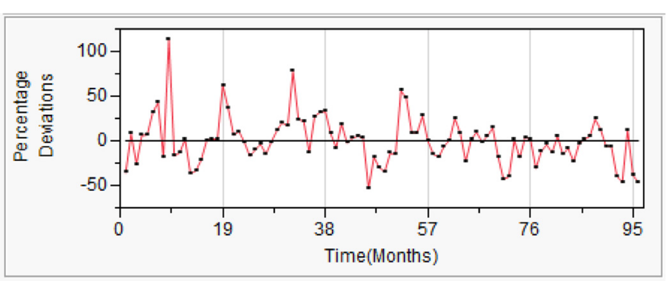

Fig. A.1. Percentage deviations from the available inventory mean for all eight fiscal years.

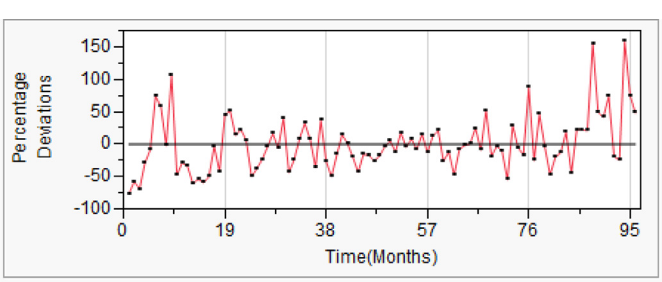

Fig. A.2. Percentage deviations from the donation mean over the eight fiscal years.

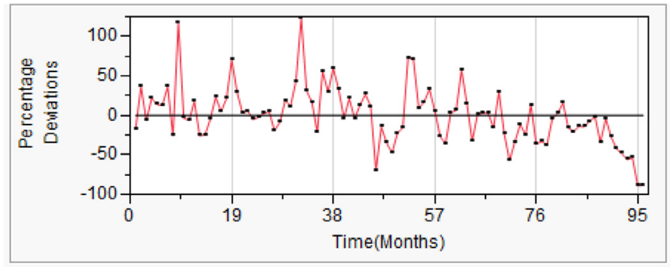

Fig. A.3. Percentage deviations from the transfer mean over the 8 fiscal years.

Table A.1
Summary of the discretized percentage deviations of each dataset.

Variable	Lower bound	Upper bound	# of bins
Available inventory	-50	90	16
Donations	-70	90	18
Transfer in	-80	90	19

A.2. Hypothesis testing of stationary transition probabilities (Anderson & Goodman, 1957)

This hypothesis test will verify our assumption that the probability of the system moving from the current state v, to the next state v', $p(v'|v)$ is stationary over the time horizon.

The alternative to this assumption is that the transition probabilities are dependent on time. The transition probabilities are calculated as follows:

$$p\left(v'|v\right) = \frac{n\left(v'|v\right)}{\sum_{v'} n(v'|v)} \tag{A.2.1}$$

$$p\left(v'|v,\ t\right) = \frac{n\left(v'|v,\ t\right)}{\sum_{v'} n(v'|v,\ t)} \tag{A.2.2}$$

Where, $n(v'|v)$ denote the number of times the system moves from current state v, to the next state v' and $n(v'|v,\ t)$ denote the number of times the system moves from current state v, to the next state v' at time t.

$$H_0 : p\left(v'|v,\ t\right) = p\left(v'|v\right)\ \forall\ t = 1,\ 2, 3, \ldots,\ 12$$

$$H_1 : p\left(v'|v,\ t\right) \neq p\left(v'|v\right)\ \forall\ t = 1,\ 2, 3, \ldots,\ 12$$

Significance level = 0.05

Decision: H_0 is rejected when the p values are smaller than 0.05.

Conclusion

At the 0.05 significance level the results in Table A.2 indicates that the transition probabilities are stationary since the p values are greater than 0.05.

A.3. Augmented Dickey–Fuller (ADF) test for stationarity

The series are stationary if their absolute test statistics (−7.36 and −5.75 for donation and transfer, respectively) are greater than the absolute critical value (2.93) (Sjö, 2008). The results are shown in Table A.3.

Table A.2
Summary of results for hypothesis testing.

State (ν)	Chi sq ($\boldsymbol{x}_\nu{}^2$)	$P(\boldsymbol{X} > \boldsymbol{x}_\nu{}^2)$	DF	Decision
1	0	1	66	Do not reject Ho
2	0	1	66	Do not reject Ho
3	5.142857143	1	66	Do not reject Ho
4	6.518518519	1	66	Do not reject Ho
5	24.39861111	0.999	66	Do not reject Ho
6	24.21607906	0.999	66	Do not reject Ho
7	33.71828704	0.999	66	Do not reject Ho
8	0	1	66	Do not reject Ho
9	26.76653439	0.999	66	Do not reject Ho
10	0	1	66	Do not reject Ho
11	0	1	66	Do not reject Ho
12	0	1	66	Do not reject Ho
13	0	1	66	Do not reject Ho
14	0	1	66	Do not reject Ho
15	0	1	66	Do not reject Ho
16	0	1	66	Do not reject Ho
Total	120.7608873	1	1056	Do not reject Ho

Table A.3
Summary of results for stationarity test for transformed data (donations and transfers).

Summary of results for stationarity test	Donation % deviation	Transfer % deviation
Mean	2.96e−16	2.08e−4
Standard deviation	42.95	36.05
N	96	96
Zero mean ADF	-7.40	-5.78
Single mean ADF	-7.36	-5.75
Trend ADF	-7.96	-7.01

A.4. Analytical results

A.4.1. Derivation of minimum percentage adjustment in sample mean

Let ν be the available inventory (in pounds) associated with pseudo-state $\tilde{\nu}$. The allocation decisions are made from available inventory and incoming donations. Given the available inventory is allocated first, the unmet need in state ν that must be met through donations is $[\sum_{c\in C} d_c - \nu]^+$. This implies the donation amount must satisfy the following condition.

Condition 1: $X \geq [\sum_{c\in C} d_c - \nu]^+$.

Denote the lowest possible realization for the random variable X as $\mu_d(1+\tilde{x}_{LB})$, where $\tilde{x}_{LB}$ is the lower bound for the pseudo state values. Then to determine the percentage adjustment needed in the donation sample mean that satisfies the condition 1 we must have $(1+\alpha_\nu)\ \mu_d(1+\tilde{x}_{LB}) \geq [\sum_{c\in C} d_c - \nu]^+$ where α_ν is the relative increase/decrease in the mean.

Solving for α_ν gives the desired expression

$$\alpha_\nu = \frac{(\sum_{c\in C} d_c - \nu)}{\mu_D * (1+\tilde{x}_{LB})} - 1.$$

A.4.2. Proof that the Proportional Allocation (PA) rule is equitable

Proposition. If history is zero or all demand is satisfied, rule PA is perfectly equitable.

Proof. Let $x_1 = Previous\ 11\ months\ for\ county\ 1$

Let $x_2 = Previous\ 11\ months\ for\ county\ 2$

Let $s = available\ supply$

$$PPIP_1 = \left(x_1 + s*\frac{P_1}{P_1+P_2}\right)*\frac{1}{P_1}$$
$$PPIP_1 = \left(\frac{x_1}{P_1}\right) + \frac{s}{P_1+P_2}$$

$$PPIP_2 = \left(x_2 + s*\frac{P_2}{P_1+P_2}\right)*\frac{1}{P_2}$$

$$PPIP_2 = \left(\frac{x_2}{P_2}\right) + \frac{s}{P_1+P_2}$$

$$\text{If } \left(\frac{x_1}{P_1}\right) \text{or} \left(\frac{x_2}{P_2}\right) = 0, \quad \text{then} \quad PPIP_1 = PPIP_2$$

□

When the previous 11-month history is based on fully satisfying the county demand, then

$x_i = \frac{(P_i*75)*11}{12}$, Therefore $\frac{x_1}{P_1} = \frac{x_2}{P_2} = 75*\frac{11}{12}$ and therefore $PPIP_1 = PPIP_2$

A.4.3. Proof that the unmet need (PPIP) for rule SSDF is smaller than unmet need for PA and SLDF rule

Proposition. If the total demand exceeds the total supply, the total unmet need across all counties is larger under allocation rule PA than SSDF.

Proof.

$$Let\ U = total\ unmet\ need\ for\ all\ counties$$
$$Let\ u_i = unmet\ need\ for\ county\ i$$
$$Let\ P_i = poverty\ population\ for\ county\ i$$
$$Let\ d_i = demand\ for\ county\ i$$
$$Let\ q_i = quantity\ received\ by\ county\ i$$
$$Let\ T = Target\ PPIP$$
$$Let\ S = Total\ available\ supply$$
$$Let\ n = number\ of\ counties$$
$$u_i = T - \frac{q_i}{P_i}$$

□

Assuming $n = 2,\ d_1 < d_2\ and\ d_1 \geq S\ \rightarrow d_1 + d_2 \geq S$

For SSDF

$$q_1 = \min(S,\ d_1) = S$$
$$q_2 = \min(S - q_1,\ d_2) = \min(S - S,\ d_2) = 0$$
$$U_{SSDF} = T - \frac{S}{P_1} + T - \frac{0}{P_2} = 2T - \frac{S}{P_1}$$

For SLDF

$$q_2 = \min(S,\ d_2) = S$$
$$q_1 = \min(S - q_2,\ d_1) = \min(S - S,\ d_2) = 0$$
$$U_{SLDF} = T - \frac{0}{P_1} + T - \frac{S}{P_2} = 2T - \frac{S}{P_2}$$

For PA

$$q_i = \min\left(S*\frac{P_i}{\sum_i P_i},\ d_i\right)$$
$$u_i = T - S*\frac{P_i}{P_i\sum_i P_i}$$
$$u_i = T - \frac{S}{\sum_i P_i}$$
$$U = \sum_i T - \frac{S}{\sum_i P_i}$$
$$U_{PA} = 2T - \frac{2S}{P_1+P_2}$$
$$d_i \propto P_i,\ since\ d_1 < d_2 \rightarrow\ P_1 < P_2$$

Table A.4
Expected unmet need (PPIP) for all counties.

Pseudo state (%)	State	Expected unmet need (PPIP) for all counties SSDF	SLDF	PA
<-50	1	44	214	131
-45	2	43	213	130
-35	3	42	208	126
-25	4	41	202	124
−15	5	39	195	120
-5	6	38	189	117
5	7	37	182	112
15	8	36	176	108
25	9	34	171	105
35	10	33	165	101
45	11	32	159	98
55	12	31	153	94
65	13	30	147	90
75	14	28	141	87
85	15	27	135	83
>90	16	27	132	82

Comparing PPIP;

$$\frac{S}{P_2} < \frac{S}{P_1}$$

$$\frac{S}{2P_2} < \frac{S}{P_1 + P_2} \rightarrow \frac{S}{P_2} < \frac{2S}{P_1 + P_2}$$

Comparing Unmet need;

$$U_{SSDF} = 2T - \frac{S}{P_1}$$

$$U_{SLDF} = 2T - \frac{S}{P_2}$$

$$U_{PA} = 2T - \frac{2S}{P_1 + P_2}$$

$$U_{SSDF} < U_{PA} < U_{SLDF}$$

Discussion. We provide a brief discussion of the results for the finite horizon model under the base scenario. The results from Table A.4 show that unsatisfied demand per person in poverty (determined as the sum of the unsatisfied demand from all counties) is smaller for allocation rule SSDF compared to SLDF and PA. For allocation rule SLDF, most of the counties are underserved except the county with the largest demand (Durham County which contributes to about 44% of the entire poverty population being served) resulting in most counties being under served. More counties being under served constitutes to higher unsatisfied demand. On the other hand, in the SSDF allocation rule, most of the counties are well served except the county with the largest demand hence reducing the unsatisfied demand per person in poverty.

References

Ahiska, S. S., Appaji, S. R., King, R. E., & Warsing, D. P., Jr (2013). A Markov decision process-based policy characterization approach for a stochastic inventory control problem with unreliable sourcing. *International Journal of Production Economics, 144*(2), 485–496. http://dx.doi.org/10.1016/j.ijpe.2013.03.021.

Altay, N., & Green III, W. G. (2006). OR/MS research in disaster operations management. *European Journal of Operational Research, 175*(1), 475–493. http://dx.doi.org/10.1016/j.ejor.2005.05.016.

Anderson, T. W., & Goodman, Leo A. (1957). Statistical Inference about Markov Chains. *Annals of Mathematical Statistics, 28*(1), 89–110. doi:10.1214/aoms/1177707039.

Balcik, B., Beamon, B. M., Krejci, C. C., Muramatsu, K. M., & Ramirez, M. (2010). Coordination in humanitarian relief chains: Practices, challenges and opportunities. *International Journal of Production Economics, 126*(1), 22–34. http://dx.doi.org/10.1016/j.ijpe.2009.09.008.

Balcik, B., & Beamon, B. M. (2008). Facility location in humanitarian relief. *International Journal of Logistics Research and Applications, 11*(2), 101–121. doi:10.1080/13675560701561789.

Beamon, B. M., & Kotleba, S. A. (2006). Inventory modelling for complex emergencies in humanitarian relief operations. *International Journal of Logistics Research and Applications, 9*(1), 1–18. doi:10.1080/13675560500453667.

Begen, M. A., Pun, H., & Yan, X. (2016). Supply and demand uncertainty reduction efforts and cost comparison. *International Journal of Production Economics, 180*, 125–134. http://dx.doi.org/10.1016/j.ijpe.2016.07.013.

Blanco, E. E., & Goentzel, J. (2006). *Humanitarian supply chains: A review*. MIT Center for Transportation & Logistics, POMS.

Bozorgi-Amiri, A., Jabalameli, M. S., & Mirzapour Al-e-Hashem, S. M. J. (2013). A multi-objective robust stochastic programming model for disaster relief logistics under uncertainty. *OR Spectrum, 35*(4), 905–933. doi:10.1007/s00291-011-0268-x.

Chomilier, B., Samii, R., & Van Wassenhove, L. N. (2003). The central role of supply chain management at IFRC. Forced Migration Review. (18), 15–16. Retrieved from doi:http://www.fmreview.org/en/FMRpdfs/FMR18/fmr1807.pdf. Accessed 1.29.2016.

Coleman-Jensen, A., Rabbitt, M. P., Gregory, C., & Singh, A. (2015). *Household food security in the United States in 2014. ERR-194, U.S. Department of Agriculture, Economic Research Service.* (pp. 6–10).

Davis, L. B., Samanlioglu, F., Qu, X., & Root, S. (2013). Inventory planning and coordination in disaster relief efforts. *International Journal of Production Economics, 141*(2), 561–573. http://dx.doi.org/10.1016/j.ijpe.2012.09.012.

FeedingAmerica. (2014). About Us. Retrieved from doi: http://feedingamerica.org/how-we-fight-hunger/about-us.aspx.

Food Bank of Central and Eastern North Carolina. (2012). *Fairshare calculations*. Food Bank of Central and Eastern North Carolina.

Food Bank of Central and Eastern North Carolina. (2016). About Us. Retrieved from http://www.foodbankcenc.org/site/PageServer?pagename=about. Accessed 1.4.2017.

Galindo, G., & Batta, R. (2013). Review of recent developments in OR/MS research in disaster operations management. *European Journal of Operational Research, 230*(2), 201–211. http://dx.doi.org/10.1016/j.ejor.2013.01.039.

Güllü, R., Önol, E., & Erkip, N. (1999). Analysis of an inventory system under supply uncertainty. *International Journal of Production Economics, 59*(1–3), 377–385. http://dx.doi.org/10.1016/S0925-5273(98)00024-3.

Howard, R. A. (1960). *Dynamic programming and Markov processes*. Cambridge, MA: MIT Press.

Karaesmen, I., Scheller-Wolf, A., & Deniz, B. (2010). Managing perishable and aging inventories: Review and future research directions. In K. G. Kempf, P. Keskinocak, & R. Uzsoy (Eds.), *Planning production and inventories in the extended enterprise: A state of the art handbook,: 1* (pp. 393–436). New York: Springer.

Lien, R. W., Iravani, S. M. R., & Smilowitz, K. R. (2014). Sequential resource allocation for nonprofit operations. *Operations Research, 62*(2), 301–317. doi:10.1287/opre.2013.1244.

Marsh, M. T., & Schilling, D. A. (1994). Equity measurement in facility location analysis: A review and framework. *European Journal of Operational Research, 74*(1), 1–17. http://dx.doi.org/10.1016/0377-2217(94)90200-3.

Mohan, S., Gopalakrishnan, M., & Mizzi, P. J. (2013). Improving the efficiency of a non-profit supply chain for the food insecure. *International Journal of Production Economics, 143*(2), 248–255. http://dx.doi.org/10.1016/j.ijpe.2011.05.019.

Mohebbi, E. (2003). Supply interruptions in a lost-sales inventory system with random lead time. *Computers & Operations Research, 30*(3), 411–426. http://dx.doi.org/10.1016/S0305-0548(01)00108-3.

Mohebbi, E. (2004). A replenishment model for the supply-uncertainty problem. *International Journal of Production Economics, 87*(1), 25–37. http://dx.doi.org/10.1016/S0925-5273(03)00098-7.

Orgut, I. S., Ivy, J., Uzsoy, R., & Wilson, J. R. (2015). Modeling for the equitable and effective distribution of donated food under capacity constraints. *IIE Transactions, 48*(3), 252–266.

Orgut, I. S., Brock III, L. G., Davis, L. B., Ivy, J. S., Jiang, S., Morgan, S. D., et al. (2016). Achieving equity, effectiveness, and efficiency in food bank operations: strategies for feeding America with implications for global hunger relief. *Advances in managing humanitarian operations* (pp. 229–256). Chicago: Springer International Publishing.

Ozbay, K., & Ozguven, E. (2007). Stochastic humanitarian inventory control model for disaster planning. *Transportation Research Record: Journal of the Transportation Research Board, 2022*(1), 63–75. doi:10.3141/2022-08.

Puterman, M. L. (2009). *Markov decision processes: Discrete stochastic dynamic programming*. Wiley.

Rottkemper, B., Fischer, K., & Blecken, A. (2012). A transshipment model for distribution and inventory relocation under uncertainty in humanitarian operations. *Socio-Economic Planning Sciences, 46*(1), 98–109. http://dx.doi.org/10.1016/j.seps.2011.09.003.

Sjö, B. (2008). Testing for Unit Roots and Cointegration. Lectures in Modern Econometric Time Series Analysis. Retrieved from https://www.iei.liu.se/nek/ekonometrisk-teori-7-5-hp-730a07/labbar/1.233753/dfdistab7b.pdf. Accessed 6.6.2016.

Silbermayr, L., & Minner, S. (2014). A multiple sourcing inventory model under disruption risk. *International Journal of Production Economics, 149*(0), 37–46. http://dx.doi.org/10.1016/j.ijpe.2013.03.025.

Tajbakhsh, M., Zolfaghari, S., & Lee, C. (2007). Supply uncertainty and diversification: A review. In H. Jung, F. Chen, & B. Jeong (Eds.), *Trends in supply chain design and management* (pp. 350–368). London: Springer-Verlag.

Tomlin, B. (2006). On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science, 52*(5), 639–657. doi:10.2307/20110543.

Tzeng, G.-H., Cheng, H.-J., & Huang, T. D. (2007). Multi-objective optimal planning for designing relief delivery systems. *Transportation Research Part E: Logistics and Transportation Review, 43*(6), 673–686. http://dx.doi.org/10.1016/j.tre.2006.10.012.

United States Department of Agriculture Economic Research Service (2006). *Definitions of food security*. United States Department of Agriculture Economic Research Service. Retrieved from http://www.ers.usda.gov/topics/food-nutrition-assistance/food-security-in-the-us/definitions-of-food-security.aspx Accessed 6.3.2016.

Van Wassenhove, L. N. (2005). Humanitarian aid logistics: Supply chain management in high gear[dagger]. *The Journal of the Operational Research Society, 57*(5), 475–489.

Whybark, C. , D. (2007). Issues in managing disaster relief inventories. *International Journal of Production Economics, 108*(1–2), 228–235. http://dx.doi.org/10.1016/j.ijpe.2006.12.012.

Wight, V., Kaushal, N., Waldfogel, J., & Garfinkel, I. (2014). Understanding the link between poverty and food insecurity among children: Does the definition of poverty matter. *Journal of Children and Poverty, 20*(1), 1–20. doi:10.1080/10796126.2014.891973.

Yang, Z. B., Aydin, G., Babich, V., & Beil, D. R. (2012). Using a dual-sourcing option in the presence of asymmetric information about supplier reliability: Competition vs. diversification. *Manufacturing & Service Operations Management, 14*(2), 202–217.

Yano, C. A., & Lee, H. L. (1995). Lot sizing with random yields: A review. *Operations Research, 43*(2), 311–334.

Yeo, W. M., & Yuan, X.-M. (2011). Optimal inventory policy with supply uncertainty and demand cancellation. *European Journal of Operational Research, 211*(1), 26–34. http://dx.doi.org/10.1016/j.ejor.2010.10.031.