Supply Management for Rapid-Onset Disasters under Demand, Supply, and Budget Uncertainty

Mahyar Eftekhar
W.P. Carey School of Business, Arizona State University

Jing-Sheng Song
Fuqua School of Business, Duke University

Scott Webster (scott.webster@asu.edu)
W.P. Carey School of Business, Arizona State University

Abstract

We study relief item supply management for the immediate relief period of rapid onset disasters. We consider two common alternatives—pre-positioning and local purchasing—and provide insights and methods for pre-positioned inventory targets. With sufficient budget, inventory should increase in disaster frequency and severity; the reverse is true otherwise. If the budget is limiting investment in inventory, then the rate of savings from improved forecasting is amplified for critical items and attenuated for noncritical items. Our model can be used to estimate the value of mitigating constraints on local spend, such as a line of credit underwritten by a donor.

Keywords: Emergency relief operations, supply management policies

Introduction

This paper investigates supply management for the *immediate relief period* of a rapid-onset disaster during which humanitarian organizations (HOs) supply relief items to beneficiaries. The immediate relief period is the first stage of an emergency response where life-savings efforts are the primary focus, and is succeeded by a *maintenance and control* stage where the situation stabilizes, and a *recovery* stage. The duration of the immediate relief period depends on many factors, including the disaster type, magnitude, and location. More predictable rapid-onset disasters tend to have shorter durations (e.g., 24 hours for hurricanes and military conflicts) compared to less predictable events (e.g., five days for earthquakes and tornadoes).

The immediate relief period is the most crucial time window, and during this time, the local infrastructure may be damaged and the abilities of local actors might be severely restricted. When a disaster strikes, HOs estimate demand (i.e., what items are required and in what quantity), transport the relief items (either from HO warehouses or from a local supplier) to the field, and distribute among the beneficiaries. Due to many unknowns such as time, place, type, and magnitude of a disaster, supply management is a significant challenge. The limited financial resources of HOs combined with the life-critical nature of the immediate relief period demands that each dollar spent on supply management be used as effectively as possible.

We study two common supply management alternatives—pre-positioning and local purchasing—and analyze policies that minimize the expected sum of purchase, transport, storage, and shortage costs. We consider the setting where locally purchased items are prioritized over pre-positioned stock, which is common in practice for several reasons: (1) locally purchased items are often less expensive due to lower transportation costs, (2) are more culturally acceptable due to regular usage among the affected population, and (3) purchases from local suppliers helps speed the economic recovery of the affected area (see Table 1).

Table 1 – Comparison between prepo stock and reactive stock

Aspect	Prepo Stock	Reactive Stock
Demand	Unknown	Known
Supply	Reliable	Unreliable
Purchase price	Stable / low	Unstable / high
Holding cost	Yes	No / ignorable
Transportation cost	Very high	Low
Other advantages	Product quality	Culturally accepted products Stimulation of local economy

Literature Review

Supply management in emergency relief operations has drawn significant attention in the Operations Management literature, most of which focuses on the proactive approach. Topics include optimal pre-positioning locations (Balcik and Beamon 2008, de Treville et al. 2006, Duran et al. 2011, Manoj et al. 2016), response capacity (Beamon and Kotleba 2006, Kunz et al. 2014, Salmeron and Apte 2010), and location-allocation models (Mete Zabinsky 2010). These studies typically analyse a two-stage linear stochastic program for optimizing prepo stock of a single item (Acimovic and Goentzel 2016, Manoj et al. 2016, Rawls and Turnquist 2012, Salmeron and Apte 2010). While this body of the literature provides a comprehensive picture of the proactive approach, it ignores the possibility of local purchase, and the consequent tradeoff between prepo and reactive stocks. Neither the supply uncertainty nor the budget uncertainty has been considered in these works.

Demand uncertainty and supply chain disruption have long been studied in OM literature. Similar to the proactive approach in humanitarian settings, a common strategy of commercial firms to hedge against supply disruption and random demand surges is to maintain stockpiles (Liu et al. 2016, Sheffi 2005) whose size and source of supply are determined based on a cost-benefit analysis. For instance, Tomlin (2006) considers a case where a firm can choose to buy either from a reliable but expensive supplier, or unreliable but cheaper supplier, and shows that a risk neutral firm will pursue a pure disruption-management strategy by carrying inventory, and single sourcing from the reliable supplier. Liu et al. (2016) explain that stockpiles are not beneficial until a supply disruption or demand surge occurs. Therefore, considering a fast-moving commodity such as pharmaceutical products, they suggest a policy that allows virtual stockpile pooling to optimize the use of stockpiles and minimize the overall inventory holding costs. Perhaps the most relevant reference to the present work is Huang et al. (2016) who study joint reactive capacity and safety stock policies where demand suddenly surges. Considering several aspects of demand surges such as duration, intensity, compactness, volatility, and frequency, they propose a policy to minimize the long-term average expenditures under a fixed service level. Our paper has fundamental differences from Huang et al. (2016). In particular, we consider supply uncertainty and budget constraints and we consider the case where reactive stock is prioritized over prepo stock (instead of vice-versa).

This paper contributes to the literature in multiple aspects. First, while most academic studies consider either supply or demand uncertainty, we consider a situation where both types of uncertainty present. Second, we introduce a new source of uncertainty, budget uncertainty, specific to humanitarian settings which has been largely ignored. Furthermore, except for Duran et al. (2011), a critical assumption in the literature is that HOs' preference is to either supply relief items only from prepo stock, or to use prepo stock first and use reactive stock only if prepo stock is insufficient to cover demand. However, the HOs' executives whom we interviewed emphasized the opposite direction; HOs' preference is to fulfill the emergency demand during immediate relief period from the local markets first, and then cover the additional demands from prepo stock.

Model and Interpretations of Results from Analysis

We briefly describe the processes and inputs used for setting prepo stock levels in practice. At the end of an immediate relief period, HO management updates its target levels of prepo stock in preparation for the next disaster. The underlying assumption is that prepo stock used during a disaster will be replenished before the next disaster. One key input for setting target prepo stock levels for different products is an estimate of demand during the immediate relief period. Management considers past experience to establish a range on the number of individuals affected and the corresponding product demand per person. This information, in conjunction with historical experience on local supply quantities, is used to arrive at a target quantity of prepo stock. This target quantity is translated into dollars considering both purchase price and transport cost. The funds for investment in prepo stock come from within the organization (*private funds*), as opposed to appealing to donors. Management determines the target prepo stock investment considering available funds and competing priorities for these funds.

We model a single item (e.g., a kit containing a collection of critical items for survival) with prepo stock level as a decision variable that is to be determined at the end of the previous immediate relief period. This prepo stock is to be deployed in the next immediate relief period, if the reactive stock procured during that period is insufficient to cover the demand (see Figure 1).

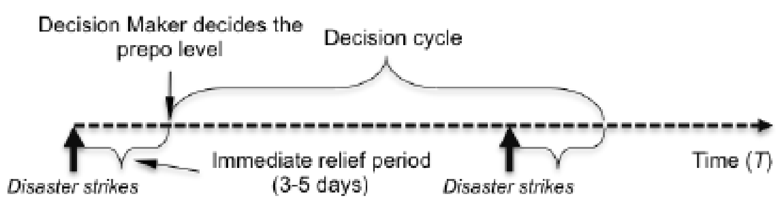

Figure 1 – Decision cycle

Next, we describe our model of the problem and provide interpretations of propositions that characterize optimal prepo and comparative statics under different conditions. Our model description serves to clarify the foundation that underlies our conclusions. To save space, we do not include the propositions (please contact an author for an expanded paper with six propositions including derivations and proofs). Elements of our model are listed below.

D = uncertain demand during immediate relief period
Q = uncertain local supply during immediate relief period, possibly correlated with D
T = uncertain time to next disaster, which is independent of D and Q
μ_j, σ_j = mean and standard deviation of $j \in \{D, Q, T\}$
c = cost per unit (including transport) for prepo stock, normalized to 1, i.e., $c = 1$
α = ratio of local cost to prepo cost per unit (including transport), or *local supply cost multiple*, $\alpha < 1$
i = cost per dollar-period for holding prepo stock
v = cost per unit of unsatisfied demand during the immediate relief period, $v > c$
b = available budget (including current prepo inventory) at the start of the cycle
r = inflow per period during a cycle (i.e., used for prepo inventory holding cost and purchase of local stock during the immediate relief period)
x = order-up-to quantity of prepo stock at start of the cycle, with $x \le r/i$ (i.e., prepo holding cost per period must not be more than cash inflow rate)
$S(x) = \left(D - \min\left\{ Q, \frac{b + rT - (1+iT)x}{\alpha} \right\} \right)^+$ = uncertain *local shortage*, i.e., extent to which demand exceeds the local purchase quantity where local purchase quantity is determined from the smaller of available local supply and available funds
$\left(S(x) - x\right)^+$ = uncertain *global shortage*, i.e., extent to which demand exceeds the total of local purchase quantity and prepo

The expected cost during a cycle and the decision problem are

$$C(x) = E\left[iTx + \alpha \min\left\{ D, Q, \frac{b + rT - (1+iT)x}{\alpha} \right\} + \min\left\{ x, S(x) \right\} + v\left(S(x) - x\right)^+ \right] \tag{1}$$

$$\Pi: \min_{x \in [0, \min\{b, r/i\}]} C(x). \tag{2}$$

We identify an expression that gives a threshold funding level that delineates the complexity of problem Π. The threshold may be expressed in terms of either cash inflow rate (denoted r_τ) or initial budget (denoted b_τ).

Interpretations of Results when Funding is more than the Threshold
If funding is above the threshold, then we obtain a relatively simple expression for optimal prepo, which balances the cost of holding inventory against the cost of a prepo stock shortage in the event that the local purchase quantity is insufficient to cover demand. From comparative statics analysis, we offer lessons for how management may adjust prepo stock in response to changes in the environment. Some of the directional results follow from an understanding of the tradeoff discussed above. For example, from this tradeoff, it follows rather naturally that increases in holding and purchase cost put downward pressure on prepo stock whereas increases in disaster frequency, shortage cost, and funds put upward pressure on prepo stock. And, as local stock is prioritized over prepo stock, changes in the local supply cost multiple have no effect on prepo stock.

The directional effects related to moments and correlation of random variables are not directly tied to the above tradeoff, and warrant discussion. First, changes in the volatility of time between disasters (σ_T) have no effect on prepo stock. Funding is sufficient to assure that local purchases will not be limited by the budget. As a consequence, the uncertain funding available for local purchases does not enter into the prepo stock tradeoff. Of course, optimal prepo stock is sensitive to changes in the first moment of T (μ_T) through its effect on holding cost. Second, as one may expect, optimal prepo increases as expected demand increases and as expected local supply decreases.

The impacts of changes in uncertainty and correlation of demand and supply are more subtle due to dependence on the shortage criticality of the item. We label a relief item as *critical* if the optimal unconstrained prepo stock exceeds the expected mismatch between demand and local supply and *noncritical* if the opposite inequality holds. As the uncertainty in demand and local supply increase, the optimal prepo stock of critical relief items increases and the optimal prepo stock of noncritical relief items decreases. An increase in uncertainty increases the tails of the random mismatch ($D – Q$) distribution, which decreases the shortage probability for a critical relief item, thus requiring an increase in prepo to compensate (and vice-versa for a noncritical item). The same behavior is associated with an increase in negative correlation between D and Q. This is because the introduction of negative correlation amplifies the variance of mismatch between demand and local supply (i.e., variance of $D – Q$), just as with increases in σ_D and/or σ_Q.

In addition to understanding whether prepo stock should increase or decrease in response to changes in the environment, the directional impacts provide insight into the impact of candidate interventions by management. For example, it is clear that investments to improve forecasts of demand and/or local supply help reduce the costs of mismatch between demand and supply. However, such investments are likely to be more attractive for critical items than noncritical items. One relatively obvious reason for this is the high shortage cost of critical items compared to noncritical items, e.g., with no change in prepo stock, the value of improved forecasting is increasing in shortage cost. However, we show that optimal prepo stock of a critical relief item decreases as forecasts improve. Thus, in settings where the budget is limiting investment in prepo, the rate of savings from improved forecasts is amplified for critical items, reflecting gains from both lower mismatch cost and lower cost of insufficient funding. In contrast, the inventory effect for noncritical relief items is reversed, possibly exacerbating the cost of budget limitations as forecast accuracy improves. Similarly, management may consider investments that dampen the negative effects of rapid-onset disaster on local supply, which lead to increases in average supply and reduce the negative correlation between demand and supply (e.g., investments that reduce the chance that a disaster reduces or eliminates local supply). Such investments put strong downward pressure on prepo stock for critical items.

Interpretations of Results when Funding is less than the Threshold

The previous setting tends to fit HOs that provide relief resources well beyond the immediate relief period in response to a disaster, or more generally, HO’s where local purchases during the immediate relief period are not constrained by the budget. As a consequence, optimal prepo balances the cost of holding prepo inventory against the shortage cost that occurs when demand is greater than the sum of local supply and prepo stock (to the extent that the prepo funding limit of $\min\{b, r/i\}$ allows).

In contrast, this setting where funding is below the threshold tends to fit HOs that attend only to life-critical relief (e.g., exiting the region shortly after the immediate relief period), or more generally, where the budget restricts what can be purchased locally. This funding limitation introduces an additional tradeoff in the determination of optimal prepo: balancing the cost of an excess local budget against the cost of an insufficient local budget. The former is associated with a global shortage and local budget more than local supply; cost would have been lower if prepo stock was higher. The latter is associated with a local shortage and local budget less than local supply; cost would have been lower if prepo stock was lower (i.e., because it is more cost effective to cover shortages with local stock).

The shift from a simple two-dimensional tradeoff to four dimensions leads to a much more complex expression for optimal prepo and changes some of the directional effects that hold in the simpler setting. For example, increases in disaster frequency and average demand may lead to decreases (instead of increases) in optimal prepo, especially under high budget pressure (e.g., reduce prepo to save dollars for local purchase, which is less expensive).

The results discussed above offer insight into how a general trend of increasing frequency and severity of disasters (e.g., due to a combination of accelerating climate change and population growth) will affect HO supply management. If budget pressure during the immediate relief period is relatively light (i.e., likely that there will be sufficient funds for purchasing local supply), then management should increase the investment in prepo stock as disaster frequency and severity increase. If budget pressure is tight, then management should decrease the investment in prepo stock as disaster frequency and severity increase. Thus, as disasters become more frequent and severe with a consequent increase in budget pressure, we can expect HOs to place greater emphasis on local supply during the immediate relief period.

Numerical illustrations

Our goal in this section is to illustrate how an HO manager may develop reasonable prepo targets under limited information and without reliance on specialized optimization software. The expressions for lower and upper bounds on optimal prepo that we derive are relatively straightforward to compute for candidate parameter values and probability distributions. Through numerical analysis in this section, we seek to provide insight into (1) settings where optimal prepo is likely to be closer to the lower bound, closer to the upper bound, or near the middle, and (2) conditions under which expected cost is relatively insensitive to different levels of prepo between the lower and upper bound. Such insight can be useful for understanding how alternative assumptions on costs and uncertainty translate into reasonable prepo targets, and ultimately for helping to make prepo decisions that take advantage of (likely limited) available information.

In the following two paragraphs we provide an example of how probability distributions of demand, local supply, and time between relief events may be estimated when limited information is available. In general, a manager may consider a number of plausible probability distributions or scenarios, and may investigate how prepo targets change as assumptions vary. Such investigation allows a manager to identify prepo targets that are reasonably satisfactory over a range of plausible assumptions. While we present a single set of distributions in the following illustrative example, it will be clear that the process may yield a number of plausible probability distributions for analysis. We use the data from the example in our numerical illustrations.

A manager begins with a review of past events, and from this, she believes that demand during the immediate relief period of the next event should exceed local supply by no more than 50 units and will be no more than 20 units less than local supply. Demand alone should be no more than 50, and no less than 10, which in conjunction with the previous information, implies that supply could be as low as 0 (leading to $D - Q = 50$ when D is at its highest value of 50) and no more than 30 (leading to $D - Q = -20$ when D is at its lowest value of 10). Without information to suggest otherwise, she estimates that each realization of demand or supply is equally likely within its range (i.e., the marginal distributions of *D* and *Q* are uniform). She suspects that there is negative correlation between demand and supply (e.g., a large disaster leads to high demand and may destroy local supplies). With little basis to estimate the degree of negative correlation, she considers the extremes of independent *D* and *Q*, and perfect negative correlation.

Finally, she believes events occur at an average rate of about one per year, and she cannot discount the possibility of immediate relief activities for multiple events occurring nearly simultaneously. Without information to indicate otherwise, she estimates that the time between relief events is an exponential random variable. Empirically, a variety of arrival processes show evidence of exponentially distributed inter-arrival times, e.g., outbreak of wars, tornados (Hayes 2002, Richardson 1956). Theoretical support for this observation comes from the Khintchine limit theorem: under fairly mild assumptions, the distribution of time between arrivals for a process that is a superposition of *n* independent arrival processes approaches exponential as *n* increases (Feller 1965, Khintchine 1960).

Figure 2 reports optimal prepo, its lower and upper bounds, and expected costs per cycle for the cases of independent demand/supply (left plots) and demand/supply with perfect negative correlation (right plots) as budget varies. (We use stochastic optimization with 50,000 trials per simulation via Analytic Solver Platform from Frontline Systems to generate the figures.) Results are presented for a single relief item with low (v = 1.5) and high (v = 6) shortage cost. The slope of the optimal expected cost curve shows the reduction in expected cost per unit increase in budget.

We offer three main observations from the figure. First, the basic patterns in prepo curves and expected cost curves are relatively stable across changes in correlation and shortage cost rate. We see diminishing return to increases in budget and higher marginal value of budget for the high shortage cost item, both of which are to be expected. Second, the expected cost when prepo is set to the lower bound is generally closer to the optimum than when prepo is set to the upper bound. While the complexity of the model inhibits characterization of the prevalence of this phenomenon, we have found that it holds over other parameter values that we tested. Third, we see that convergence of lower and upper bounds on prepo is occurring well before the threshold b_τ where such convergence is assured (e.g., b_τ is 50 or more in the examples with convergence occurring at b = 20 or less). While we have found numerical examples where convergence occurs close to the threshold, our numerical results indicate that the bounds provide assurance of optimal prepo in some settings where the budget is well below the threshold.

Finally, we illustrate how our methods can be used to gauge when the elimination of the local spend constraint will, and will not, have a large impact on alleviating human suffering. This exercise can help HOs that are tightly constrained on local spend to evaluate whether initiatives that mitigate such a constraint (e.g., line of credit underwritten by large donors that is available during the immediate relief period) would be worth pursuing. Figure 3 reports the reduction in optimal expected cost when a constraint on local spending is relaxed. As in Figure 2, the results for a less critical (v = 1.5) and a more critical item (v = 6) are reported as the budget available for prepo ranges from 1 to 25.

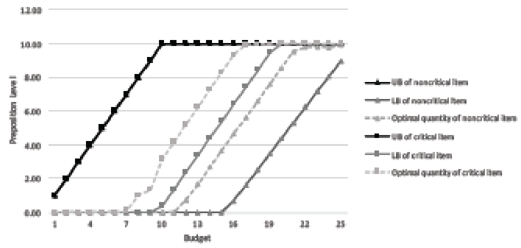

(a) Perfectly negatively correlated $D-Q$

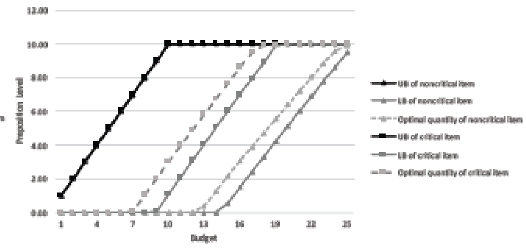

(b) Independent $D-Q$

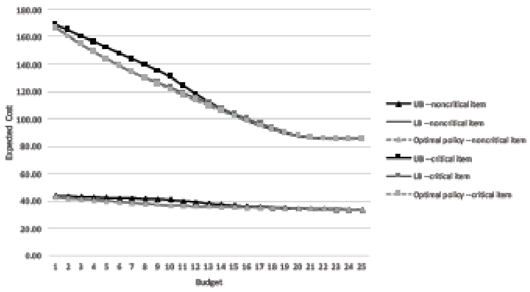

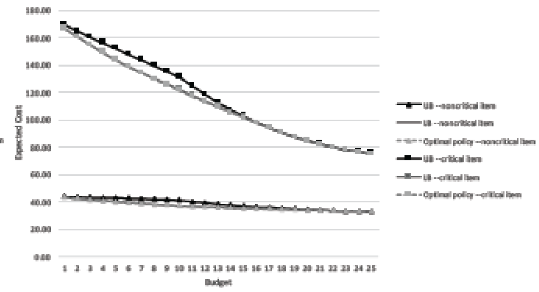

(c) Perfectly negatively correlated $D-Q$ (d) Independent $D-Q$

Figure 2 – The upper plots show the optimal prepo and lower and upper bounds under two scenarios: (1) noncritical item (v = 1.5), (2) critical item (v = 6). The lower plots show the corresponding expected cost. Parameter values: α = 0.75, i = 0.1, μ_T = 1, $v \in$ {1.5, 6}, $b \in$ [1, 25], r = 1, D ~ U(10, 50), Q ~ U(0, 30). For the left plots, b_τ = 58.5 for v = 1.5 and b_τ = 71.1 for v = 6; for the right plots, b_τ = 50.6 for v = 1.5 and b_τ = 65.6 for v = 6.

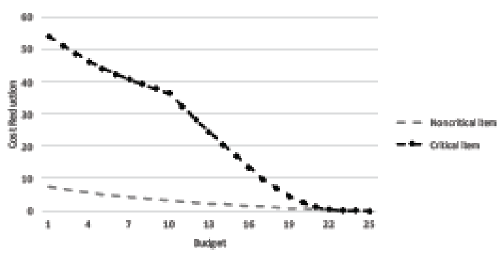

(a) Perfectly negatively correlated $D-Q$

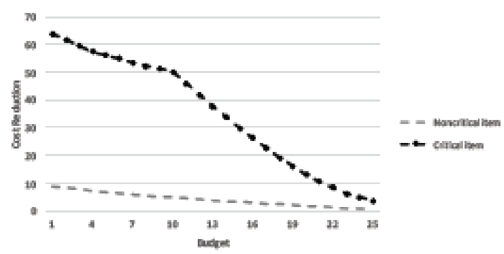

(b) Independent $D-Q$

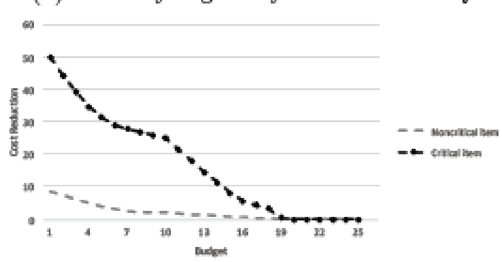

(c) Perfectly negatively correlated $D-Q$

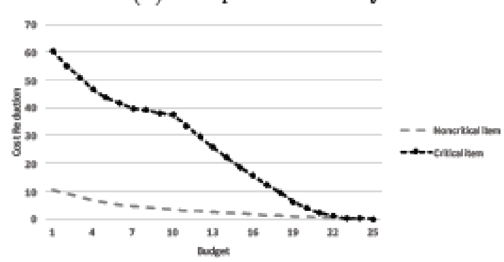

(d) Independent $D-Q$

Figure 3 – Reduction in optimal expected cost if the local spend constraint is relaxed. Parameter values: $\alpha \in$ {0.5, 0.75}, i = 0.1, μ_T = 1, $v \in$ {1.5, 6}, $b \in$ [1, 25], r = 1, D ~ U(10, 50), Q ~ U(0, 30). The top two graphs report results for α = 0.75, and bottom two graphs report results for α = 0.5.

Summary and Conclusion

Through extensive interactions with HOs and their executives, we identified several important features in HOs' practices and challenges that have not been explored in the Operations Management literature on relief item supply management. These features include the distinction between local supply (reactive stock) and central supply (prepo stock), the high transportation cost of prepo stock, the importance and priority of reactive stock, the uncertainty of both demand and local supply, budget constraint, and inflow of funds such as donations during each decision cycle. To fill this gap between theory and practice, in this paper we develop models that explicitly take into account these new features. We derive closed-form solutions and efficient algorithms to determine the optimal pre-positioned stock level in anticipation of the next disaster event, with the consideration of uncertainty of time to the next event and the associated holding cost, uncertainty of demand, priority of using (uncertain) local supply, and uncertainty of budget. We also conduct and discuss extensive comparative statics analysis to reveal insights and use numerical experiments to illustrate various effects.

The main lessons from our models and analysis for HO managers can be summarized as follows. First, we provide simple methods that a manager may use to identify ranges of reasonable prepo values under differing assumptions of cost rates and probability distributions of demand, local supply, and time between disasters (e.g., via lower and upper bounds). And our comparative statics results help a manager to gauge the directional effects of changes in parameter estimates. As a related but more speculative finding, our numerical results hint that it is less costly to error on the side of too little prepo than too much prepo. In particular, we find that even when “true”' optimal prepo is midway between lower and upper bounds, the cost at the lower bound is closer to the optimal cost (e.g., cost function is flatter to the left of optimal than to the right). In the face of limited data, a manager may wish to favor prepo values on the left side of the range of plausible values. In practice, it is often difficult to obtain accurate estimates of parameters. Our methods and results can help a manager identify a prepo target that takes a reasonable middle ground when data for estimating parameters and probability distributions are limited.

Second, initiatives to improve demand forecast accuracy and/or reduce the negative correlation between demand and local supply serve to shrink the mismatch cost between demand and supply. The higher the shortage cost rate, the larger the benefit from such investment. This conclusion is not surprising. However, we identify an interaction between such initiatives and the level of optimal prepo stock that can amplify the value for critical items and attenuate the value for noncritical items (i.e., due to loosening versus a tightening of the budget constraint). The main lesson is that a focus on mismatch cost may understate the value of such initiatives for critical items.

Third, our models allow a manager to estimate the value of relaxing a binding constraint on local spend. Large HOs may have sufficient reserves (or access to credit) that can be spent on local supply during the immediate relief period with assurance that these reserves will be replenished through donations during the event. We show that the complex four-dimensional tradeoff that arises in the presence of a binding local spend constraint dissolves into a relatively simple two-dimensional tradeoff when the constraint is relaxed. Solution ease is one practical advantage of unrestricted local spending. However, there is a more important consideration. In particular, there is a question on when the elimination of the local spend constraint will, and will not, have a large impact on alleviating human suffering. The question is meaningful because HOs that are tightly constrained on local spend may consider initiatives to mitigate such a constraint (e.g., line of credit underwritten by large donors that is available during the immediate relief period). Such an initiative requires effort on the part of HO management and a compelling case for its value. Our model can be used to make such a case through quantification of value.

To derive these results and insights, for tractability, we have made a few simplifying assumptions. For instance, we assume the key random variables have known distributions with finite supports. We also only analyze a single cycle, so our policy is myopic, although it is expected to be optimal for dynamic systems with stationary data. Finally, we assume the funding inflow rate is constant. Relaxing any of these assumptions would be a worthy future research direction. Our models and analysis here can serve as a stepping stone for these future endeavors.

Acknowledgement

We are sincerely grateful to our industry advisors, Dave A. Coddington (Senior Advisor of Catholic Relief Services' Field Supply Chain), Mary A. Jelliti (Logistics Coordinator at GOAL International), and Sarah Robbins-Penniman (Senior Advisor of Catholic Relief Services' for Supply Chain Knowledge Management), for their valuable comments and information that they shared with us throughout this project.

References

Acimovic, J., J. Goentzel (2016) "Models and metrics to assess humanitarian response capacity," *Journal of Operations Management*, Vol. 45, pp. 11-29.

Balcik, B., B. M. Beamon (2008) "Facility location in humanitarian relief," *International Journal of Logistics: Research and Applications*, Vol. 11, No. 2, pp. 101-121.

Beamon, B.M., S.A. Kotleba (2006) "Inventory modeling for complex emergencies in humanitarian relief operations," *International Journal of Logistics: Research and Applications*, Vol. 9, No. 1, pp. 1-18.

de Treville, S., I. Smith, A. Rolli, V. Arnold (2006) "Applying operations management logic and tools to save lives: A case study of the world health organization's global drug facility," *Journal of Operations Management*, Vol. 24, No. 4, pp. 397-406.

Duran, S., M. A. Gutierrez, P. Keskinocak (2011) "Pre-positioning of emergency items for CARE International," *Interfaces*, Vol. 41, No. 3, pp. 223-237.

Feller, W. (1965) *An Introduction to Probability Theory and its Applications*, Vol. II, Wiley, New York, NY.

Hayes, B. (2002) "The killing by the numbers," *Wired*, September, pp. 81-82.

Huang, L., J-S. Song, J. D. Tong (2016) "Supply chain planning for random demand surges: reactive capacity and safety stock," *Manufacturing & Service Operations Management*, in press.

Khintchine, A.Y. (1960) *Mathematical Methods in the Theory of Queueing*, Vol. II. Charles Griffin & Company, London.

Kunz, N., G. Reiner, S. Gold (2014) "Investing in disaster management capabilities versus pre-positioning inventory: A new approach to disaster preparedness," *International Journal of Production Economics*, Vol. 157, pp. 261-272.

Liu, F., J-S. Song, J. D. Tong (2016) "Building supply chain resilience through virtual stockpile pooling," *Production and Operations Management*, in press.

Manoj, U.V., S. Kumar, S. Gupta (2016) "An integrated logistics model for predictable disasters," *Production and Operations Management*, Vol. 25, No. 5, pp. 791-811.

Mete, H., Z. Zabinsky (2010) "Stochastic optimization of medical supply location and distribution in disaster management," *International Journal of Production Economics*, Vol. 126, No. 1, pp. 76-84.

Rawls, C. G., M. A. Turnquist (2012) "Pre-positioning and dynamic delivery planning for short-term response following a natural disaster," *Socio-Economic Planning Sciences*, Vol. 46, pp. 46–57.

Richardson, L. F. (1956) "Mathematics of war and foreign politics," in Newman, J.R. (Ed.), *The World of Mathematics*, Vol. 2, Prentice-Hall, Englewood Cliffs, NJ,pp. 1240–1263.

Salmeron, J., A. Apte (2010) "Stochastic optimization for natural disaster asset prepositioning," *Production and Operations Management*, Vol. 19, No. 5, pp. 561–574.

Sheffi, Y. (2005) *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage*. MIT Press, Boston, MA.

Tomlin, B. (2006) "On the value of mitigation and contingency strategies for managing supply chain disruption risks," *Management Science*, Vol. 52, No. 5, pp. 639–657.