Proceedings of the 2010 Industrial Engineering Research Conference
A. Johnson and J. Miller, eds.

Hospital Stockpiling for Disaster Planning

Abstract ID: 419

Po-Ching DeLaurentis
IU Center for Health Services and Outcomes Research
Regenstrief Institute, Inc., Indianapolis, IN 46202, USA

Elodie Adida
Department of Mechanical & Industrial Engineering
University of Illinois at Chicago, Chicago, IL 60607, USA

Mark Lawley
The Weldon School of Biomedical Engineering
Purdue University, West Lafayette, IN 47907, USA

Abstract

This paper addresses joint inventory stockpiling of medical supplies for groups of hospitals prior to a disaster. We assume patient demand is uncertain and driven by the characteristics of a variety of scenarios of a disastrous event, and hospitals have mutual aid agreements for inventory sharing in emergency. We model this problem as a noncooperative strategic game, prove the existence of a Nash equilibrium, and analyze the equilibrium solutions. We also examine a centralized model of stockpile decision making where a central decision-maker optimizes the entire system, and we compare the solutions to those of the decentralized (game) model.

Keywords
Emergency Preparedness, Game Theory, Policy Making.

1. Introduction

Public health organizations and healthcare industries have intensified their disaster preparedness and response efforts in recent years as a response to the increasing threat of terrorist attacks and natural hazards. Although government agencies usually provide general guidelines for disaster preparedness, they typically lack the detailed instructions and/or methods decision makers need to determine appropriate levels of investment in surge capacity. Surge capacity includes personnel, equipment and medical supplies which become scarce resources once a disastrous event occurs. Thus stockpiling sufficient medical supplies becomes a very important preparedness tactic. Nonetheless, government guidelines also do not address the financial implication of stockpiling supplies to healthcare providers. In addition, as Toner and Waldhorn point out, the lack of specificity and details of these guidelines have contributed to the unwillingness of hospital decision makers to commit necessary resources in preparing for adverse events [2]. In addition, the financial cost of disaster preparedness planning is one of the most critical issues that governments and hospitals need to consider. Combined with the demand uncertainties associated with an adverse event, establishing stockpile requirements for disaster response with a limited budget is very challenging [1].

Our work focuses on stockpile decision making within a community of hospitals. We adopt a game theoretical framework since hospitals commonly engage in mutual aid agreements, such as memoranda of understanding (MOU), to share supplies with those encountering shortages in emergency, and therefore the decisions made by hospitals become inter-dependent. We assume that hospitals want to meet as much of the disaster surge demand as possible, at the lowest possible cost, considering shortages incur costs. To capture the inherent uncertainty underlying disaster planning, our model incorporates various possible scenarios, each occurring with some probability and with a given demand probability distribution. We further assume that in the case of a shortage overall in the community of hospitals, each hospital is penalized proportionally to the system's overall supply shortage. This is a reasonable assumption because, firstly, hospitals' mutual aid relationships imply that supplies are shared among the hospitals involved. Secondly, patient surge demand can be re-distributed depending on hospitals' response capacities. Thirdly, if an adverse event is an infectious disease, any untreated patient (due to insufficient available surge capacity at a hospital) will further spread the disease in the community and thus create more demand to the overall health system, hence potentially contributing to further shortages at other hospitals.

This research answers the following questions: What will be the hospital stockpile decisions in a decentralized and centralized decision making settings? What are the public policy implications provided by the analytical solutions? In a decentralized setting, each hospital aims to minimize its own cost in a game theoretical framework. In a centralized setting, a coordinated decision is made by a central planner to optimize the overall costs. Our research focuses on the stockpiling of medical supplies such as personal protective equipment (PPE), which are used to protect care providers and patient families who care for the sick or injured. We explicitly exclude flu vaccines from the scope simply because vaccine production (and hence stockpiling) is not possible before the outbreak when flu strain is discovered.

2. Hospital Stockpiling Game Model

We consider a group of hospitals $\{1,\ldots,n\}$ serving a given region. Each hospital i needs to decide the stockpile level s_i of a medical item in anticipation of a disaster. In this model, the severity and type of disaster is uncertain. Thus, the overall patient demand, D, for the item in the region is also uncertain. We consider several disaster scenarios, $l = 1,\ldots,L$ with probability $q^1,\ldots,q^L$, where $\sum_{l=1}^{L} q^l = 1$. Each scenario corresponds to a certain type of disaster with a given severity. The conditional total system-wide demand under scenario l is represented by the random variable D^l (total demand in scenario l), with a cumulative distribution function, $F^l(.)$ differentiable on $\mathbb{R}$, and a probability density function, $f^l(.)$, defined on $\mathbb{R}$ and strictly positive on $[a_l, b_l] \subset \mathbb{R}^+$ (b_l may be infinite).

We impose a penalty cost on hospitals for shortages, defined by unmet patient demand. In this model, the cost per unit of shortage is p. Moreover, we assume that because of complete sharing of supply due to mutual aid agreements or complete demand redistribution in the group of hospitals, the penalty only depends on the *total* overall shortage, i.e. on the difference between the cumulative stockpile of all hospitals and the total demand realization, and not on shortages at individual hospitals. Every hospital is penalized (proportionally) when the amount of system-wide supply falls short of its realized system-wide demand. In practice, the value of the penalty cost p can be estimated by conducting a thorough cost analysis assessing all risks involved. It can also be deemed as the avoidable cost (such as probable work loss days due to an illness) should there be sufficient level of available medical supplies.

Let c_i be the stockpiling cost of one unit of supply at hospital i. Let S be the overall stockpile level of all participating hospitals such that $S = \sum_{i=1}^{n} s_i$. The total shortage, denoted U, is the difference between the total demand and the total stockpile, if this difference is non negative (it is zero otherwise). In disaster scenario, l, the total shortage amount is $(D^l - S)^+$. The total shortage penalty cost at the system's level is given by pU. Moreover, the shortage penalty cost incurred at hospital i is a fraction w_i of the total shortage cost, where $0 < w_i < 1$ and $\sum_{i=1}^{n} w_i = 1$. The fraction w_i can be interpreted as hospital i's fraction of service capacity to the overall system capacity (typically estimated by its bed size) or hospital i's share of the responsibility in the region's preparedness.

The expected system-wide shortage depends on all hospital stockpile levels $\mathbf{s} = (s_1,\ldots,s_n)$ through its dependency on S:

$$E[U] = \sum_{l=1}^{L} q^l E[U|\ l] = \sum_{l=1}^{L} q^l E[(D^l - S)^+] = \sum_{l=1}^{L} q^l \int_S^{\infty} (x - S) f^l(x) dx \qquad (1)$$

where $E[U|\ l]$ is the expected shortage amount given that scenario l occurs. The total cost at hospital i is given by $c_i s_i + p w_i U$ and also depends on all hospital stockpile levels $\mathbf{s} = (s_1,\ldots,s_n)$ through the dependency of U on S.

One important assumption of this model is $c_i < p w_i, \quad i = 1,\ldots,n$. This states that the marginal cost of stockpiling one unit at hospital i must be less than the fraction of the marginal penalty cost paid by hospital i if the system is one unit short overall. This assumption is not very restrictive as if it does not hold, no hospital would stockpile at all.

2.1 Best Response Problem

In hospital i's best response problem, the objective is to determine the stockpile level that minimizes hospital i's expected total cost, assuming that other hospital stockpile levels $s_j,\ j \neq i$ are fixed. We denote $\mathbf{s_{-i}} = \{s_j,\ j \neq i\}$. The expected total cost at hospital i is

$$J(s_i, \mathbf{s_{-i}}) = E\left[c_i s_i + p w_i U\right] = c_i s_i + p w_i E[U].$$

Note that hospital i's expected stockpiling cost, $J_i(s_i, \mathbf{s_{-i}})$, is a function of both s_i and $\mathbf{s_{-i}}$ through the term $E[U]$. Using

(1), the best response problem of hospital i can be written as

$$\min_{s_i \geq 0} J_i(s_i, \mathbf{s_{-i}}) = \min_{s_i \geq 0} \; c_i s_i + p w_i \sum_l q^l \int_S^\infty (x - S) f^l(x) dx.$$

The cost function J_i of hospital i is a continuous, twice differentiable, convex function of s_i.

Let $\bar{G}(s) = \sum_l q^l F^l(s)$ defined on $\mathbb{R}$, strictly increasing on $[a,b]$, where $a = \min_l a_l$ and $b = \max_l b_l$ with values in [0,1]. We denote $G(.)$ the restriction of $\bar{G}(.)$ to $[a,b]$, thus $G(.)$ is strictly increasing and invertible from $[a,b]$ into $[0,1]$, and $G^{-1}(.)$ is strictly increasing from $[0,1]$ into $[a,b]$. The best response function can be written as

$$s_i^*(\mathbf{s_{-i}}) = \begin{cases} 0, & \text{if } G^{-1}(1 - \frac{c_i}{pw_i}) - S_{-i} < 0 \\ G^{-1}(1 - \frac{c_i}{pw_i}) - S_{-i}, & \text{else.} \end{cases}$$

where $S_{-i} = \sum_{j \neq i} s_j = S - s_i$. In particular, $s_i^*(\mathbf{s_{-i}})$ is piecewise linear non increasing with s_j, $j \neq i$ and with S_{-i}, with slope equal to 0 or -1.

The best response function is illustrated in Figure 1 in the case of two hospitals ($n = 2$). We observe that there is

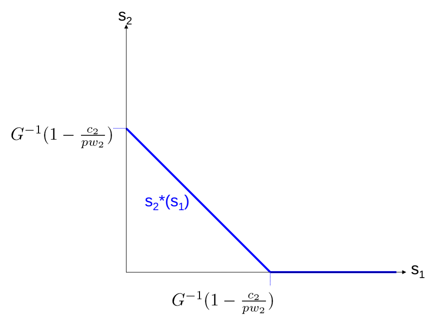

Figure 1: Best response function illustrated in the case with $n = 2$.

a threshold value $Q_i = G^{-1}(1 - \frac{c_i}{pw_i})$ which is the minimum total stockpile that hospital i desires the community to have. Thus, if all other stockpiles combined have already attained this value, hospital i stockpiles nothing additional. Otherwise, hospital i will fill the gap as needed to reach that threshold. This threshold value is such that $c_i = pw_i \Pr(D \geq Q_i)$, i.e. Q_i is the value of the total cumulative stockpile where the marginal cost at hospital i of adding one extra unit to its stockpile is equal to the expected marginal penalty paid by hospital i if this unit is not added. In other words, Q_i is the quantity in the cumulative stockpile such that hospital i is indifferent to adding one unit to its own stockpile; as long as the marginal cost of adding one extra unit is smaller than the expected penalty for not adding it, hospital i keeps adding to its stockpile.

2.2 Nash Equilibrium Solution

Inventory levels $(\bar{s_1}, \ldots, \bar{s_n})$ form a Nash equilibrium if no firm can improve its expected cost by unilaterally altering its stockpile level:

$$\bar{s_i} = s_i^*(\bar{\mathbf{s}}_{-i}), \quad i = 1, \ldots, n.$$

Graphically, a Nash equilibrium is any point that lies at the intersection of the best response functions of the players in the game. Without loss of generality, we make the assumption that the hospitals can be ordered such that

$$r_1 = \ldots = r_m < r_{m+1} \leq \ldots \leq r_n, \quad 1 \leq m \leq n,$$

where $r_i = \frac{c_i}{w_i}$. We note that since $c_i < pw_i, \quad i = 1, \ldots, n$, $0 < 1 - r_i/p (< 1)$. The Nash equilibrium solution can be expressed in closed-form as

$$\bar{s_i} = \begin{cases} \alpha_i G^{-1}(1 - \frac{r_i}{p}) & i = 1, \ldots, m \\ 0 & i = m+1, \ldots, n \end{cases}$$

for any $\alpha_1, \ldots, \alpha_m \in [0,1]$ such that $\sum_{i=1}^m \alpha_i = 1$ assuming $m > 1$, and there are infinitely many Nash equilibria. We also know that in a special case when $m = 1$, there is a unique Nash equilibrium given by

$$\overline{s_i} = \begin{cases} G^{-1}(1-\frac{r_i}{p}) & i = 1 \\ 0 & i = 2,\ldots,n. \end{cases}$$

We interpret the two results above as follows. Hospitals with a lower ratio $r_i = c_i/w_i$ have a lower per unit cost and/or receive a higher fraction w_i of the total penalty. Intuitively, having a lower ratio gives more incentives to stockpile more, as it implies that it costs less to stockpile or that more is at stake. Also, as explained above, hospital i desires the total cumulative stockpile to reach at least Q_i. Note that if $r_i = r_j$, then $Q_i = Q_j$, and a low value of r_i means a high value of Q_i. We found that the only hospitals that stockpile a positive amount at equilibrium are the hospitals with the minimum ratio r_1. These hospitals cumulatively stockpile jointly the quantity Q_1, which is the highest of all Q_i's; all others stockpile a zero quantity. When there is a single hospital with minimum ratio r_1, this hospital stockpiles its threshold quantity Q_1. When there are more than one hospital with minimum ratio r_1, there are infinitely many equilibria that correspond to different ways of splitting the total stockpile Q_1 that needs to be cumulated among all hospitals with the minimum ratio r_1. This means that hospitals with a desired quantity Q_i that is lower than Q_1 need not stockpile anything on their own as they know that others have incentive to stockpile an even greater quantity by themselves, so their own desired total quantity Q_i will be reached even if they do not contribute anything towards it, and as a result they have no incentive to stockpile at all.

Following these results, the system-wide total stockpile level at any equilibrium, S_d, is given by

$$S_d = \sum_{i=1}^{n} \overline{s_i} = G^{-1}(1-\frac{r_1}{p}), \quad \forall m, \alpha_1,\ldots,\alpha_m.$$

In particular, the total stockpile at equilibrium is the same at all Nash equilibria, even though the individual stockpiles at distinct equilibria are not the same. This total stockpile is equal to the largest of the threshold quantities $Q_i,\ i = 1,\ldots,n$, which is Q_1.

The total expected stockpiling cost for the group of hospitals at a Nash equilibrium can be written as follows:

$$\Pi_d^* = \sum_{i=1}^{n} J_i(\overline{s_i}, \overline{\mathbf{s}}_{-\mathbf{i}}) = S_d \sum_{i=1}^{m} c_i\alpha_i + p\sum_{l=1}^{L} q^l \int_{S_d}^{\infty} (x-S_d) f^l(x)dx.$$

We note in particular that Π_d^* depends on $\alpha_1,\ldots,\alpha_m$ if there are multiple equilibria; in other words the total community cost is not the same at all Nash equilibria. As a result, we define the *worst* total cost at a Nash equilibrium as the highest possible total cost over all Nash equilibria:

$$\Pi_d^{*max} = S_d c^{max} + p\sum_{l=1}^{L} q^l \int_{S_d}^{\infty} (x-S_d) f^l(x)dx$$

where $c^{max} = \max\{c_1,\ldots,c_m\}$. The worst total cost at a Nash equilibrium is obtained at the Nash equilibrium where, of all hospitals with minimum ratio $r_i = r_1$, only the hospital with maximal per unit cost stockpiles.

3. Stockpiling with Centralized Coordination

We now consider the case in which stockpile decisions at each hospital are centrally coordinated to minimize aggregate costs. One example of this setting is a situation where a healthcare network/organization or local government agency wants to determine sufficient pandemic stockpile levels of its member hospitals while minimizing the overall investment. The centralized setting can also be viewed simply as a benchmark to compare the performance of the Nash equilibrium with the best possible overall outcome. The centralized solution provides the minimal possible overall costs, should all hospitals coordinate their decisions in the interest of the entire system. The goal of the central planner is then to decide the stockpiles that minimize the expected total cost for the n hospitals:

$$\min_{\mathbf{s}=(s_1,\ldots,s_n)\geq 0} \Pi_c(\mathbf{s}) = \sum_{i=1}^{n} J_i(s_i, \mathbf{s}_{-i}) = \sum_{i=1}^{n} c_i s_i + p\sum_{l} q^l \int_{S}^{\infty} (x-S) f^l(x)dx. \tag{2}$$

The assumption of $c_i < p,\ \ i = 1,\ldots,n$ also holds here.

Let $\mathscr{I} = \{i : c_i = \min_j c_j\}$ be the set of hospitals with minimum unit stockpiling cost, and $k = |\mathscr{I}|$ the number of hospitals with minimum unit cost. We can express a closed-form solution for the centralized problem as the follows. If $k > 1$, there are infinitely many centralized solutions given by

$$\tilde{s}_i = \begin{cases} \beta_i G^{-1}(1 - \frac{c_i}{p}) & i \in \mathscr{I} \\ 0 & i \notin \mathscr{I} \end{cases}$$

for any $\{\beta_i,\ i \in \mathscr{I}\}$ such that $\beta_i \in [0,1]$ and $\sum_{i\in\mathscr{I}} \beta_i = 1$. If $k = 1$, there is a unique centralized solution given by

$$\tilde{s}_i = \begin{cases} G^{-1}(1 - \frac{c_i}{p}) & \text{if } i \in \mathscr{I} \\ 0 & \text{if } i \notin \mathscr{I}. \end{cases}$$

Since the central planner takes a system perspective and attempts to minimize the overall costs, the best strategy is to have only hospitals with the lowest per unit cost stockpile a positive quantity. If there is a single hospital with minimum unit cost, it is the only one to stockpile at the centralized optimum. If there are multiple hospitals with minimum unit cost, these hospitals may split the total quantity to stockpile among themselves arbitrarily with no effect on the total cost. The total quantity to stockpile cumulatively is the threshold value S_c such that the marginal cost of buying an extra unit of supply (at the minimum possible cost) equals the expected marginal penalty paid by the system if that unit is not purchased. In other words, $c_{i_0} = p\Pr(D \geq S_c)$ must be satisfied, where $c_{i_0} = \min_i c_i$.

The closed-form expression for the total stockpile S_c at a centralized solution follows directly from above:

$$S_c = \sum_{i=1}^{n} \tilde{s}_i = G^{-1}(1 - \frac{c_{i_0}}{p}) \quad \forall k,\ \beta_i, i \in \mathscr{I}$$

where c_{i_0} is the minimum of all the unit costs in the system, i.e., $c_{i_0} = \min_i c_i$.

In addition, the total expected stockpiling cost at a centralized solution is given by

$$\Pi_c^* = c_{i_0} S_c + p \sum_{l=1}^{L} q^l \int_{S_c}^{\infty} (x - S_c) f^l(x) dx.$$

Note that Π_c^* does not depend on $\beta_i,\ i \in \mathscr{I}$ even if there are multiple centralized solutions.

4. Comparison of the Nash Equilibrium and Centralized Solutions and Sensitivity Analysis

After examining the decentralized and centralized settings, we need to further understand the difference in their solutions, and what this implies in terms of planning preparedness and public policy. One measure of the overall performance of the solution is the aggregate total cost incurred. Another major criterion in measuring the level of preparedness is the aggregate amount of stockpile.

4.1 Efficiency of the System

We define the *worst loss of efficiency in cost* to be

$$\rho_d = \frac{\Pi_d^{*max}}{\Pi_c^*}.$$

By definition, $\rho_d \geq 1$. A ratio close to one means that the total cost in the decentralized setting is larger, but close to the total cost in the centralized setting. Thus, decentralizing or decoupling the decision-making does not lead to a large relative loss in terms of expected cost. A ratio much greater than one means the equilibrium may be very inefficient in terms of cost, when its performance is compared with that of a coordinated decision.

Let us also define the *loss of efficiency in quantity* as

$$\sigma_d = \frac{S_c}{S_d}.$$

From the definitions of S_c and S_d, we have

$$\sigma_d = \frac{G^{-1}(1 - \frac{c_{i_0}}{p})}{G^{-1}(1 - \frac{c_1}{pw_1})}.$$

The closer σ_d is to 1, the closer the total cumulative stockpile in the decentralized planning setting is to the total cumulative stockpile in the centralized setting. A larger ratio means that the equilibrium leads to a much lower total quantity than what would be achieved in a centrally coordinated setting.

Examples of the worst loss of efficiency in quantity and the loss of efficiency in cost are illustrated in Figure 2 with two hospitals. They consider $L = 3$ scenarios ($q^1 = 0.2,\ q^2 = 0.5,\ q^3 = 0.3$), and the demand distribution follows a Normal distribution with mean of 100, 300 and 500 and standard deviation of 30, 70 and 120 respectively in each of the three scenarios.

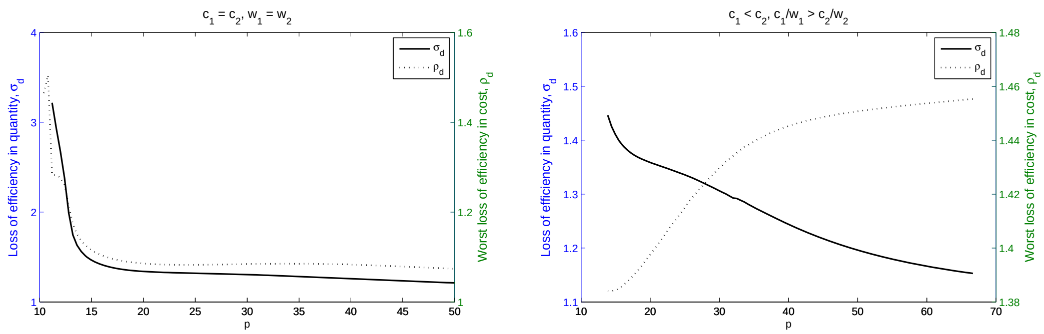

Figure 2: The loss of efficiency in quantity, σ_d , and the worst loss of efficiency in cost, ρ_d, are illustrated here for two hospitals ($n = 2$). Parameters (c_1, c_2, w_1, w_2) are given respectively: (5,5,0.5,0.5), and (4,6,0.3,0.7).

4.2 Sensitivity to the cost parameters

The effect of reducing the per unit purchasing cost on the total stockpile level is clear from the expressions of aggregate stockpile levels. As c_1 decreases, $(1 - \frac{c_1}{pw_1})$ increases leading to an increase of $S_d = G^{-1}(1 - \frac{c_1}{pw_1})$, and thus S_d increases. Similarly, S_c increases if c_{i_0} decreases. The explanation of the effect of changing the unit cost on stockpile levels is intuitive: a lower unit stockpiling cost gives incentives to stockpile more. A decrease in unit stockpiling cost may result from several different reasons. One possibility is a monetary subsidy provided by the government in the form of a subsidy per unit stockpiled as a form of cost sharing. This type of subsidy would result effectively in a decrease in the cost per unit at hospitals, and thus an increase in the total stockpile. Another situation where the unit stockpiling cost could be lowered is when the group of hospitals or health organizations in a region make their orders as one large buyer instead of several smaller buyers, resulting in a greater purchasing power for price negotiation with suppliers.

5. Conclusion

Sufficient hospital stockpile of medical supplies in preparation for a disaster is a task demanding immediate attention and action. This problem is a game theoretic model aiming at providing an estimate of how much each hospital would stockpile in a decentralized setting when minimizing its total cost. The model captures the interdependency of the decisions of hospitals serving a common population and/or participating in mutual aid agreement for sharing resources. It is compared to a centralized setting in which a coordinated stockpile decision is made for the system as a whole by a central planner while minimizing the overall expected stockpiling cost. These models provide managerial insights for public health practice in preparing sufficient medical supplies, pharmaceutical or not, to respond to adverse events.

References

[1] R. Havlak, S. E. Gorman, and S. A. Adams. Challenges associated with creating a pharmaceutical stockpile to respond to a terrorist event. *Clinical Microbiology and Infectious Diseases*, 8:529–533, 2002.

[2] E. Toner and R. Waldhorn. What hospitals should do to prepare for an influenza pandemic. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, 4(4):397–402, 2006.