Hospital Stockpiling for Influenza Pandemics with Pre-Determined Response Levels

Po-Ching C. DeLaurentis, Elodie Adida and Mark Lawley

***Abstract*— This paper explores the problem of hospital stockpiling of critical medical supplies in preparation for a possible influenza pandemic. We consider a regional network of hospitals that have mutual aid agreements in place such that they may borrow or lend supplies from each other during medical emergencies. We assume that the attack rate is a random variable with known distribution and that the demand surge due to the pandemic is a function of the attack rate, and is thus stochastic. We further assume that each hospital in the network pre-determines its targeted pandemic response level, and that any demand beyond this pre-determined level is reallocated to other hospitals. Each hospital in the network must decide the stockpile level that minimizes its expected overall cost, including purchasing cost, holding cost, cost (revenue) for borrowing (lending), penalty for setting a too low targeted level, and shortage penalty. To capture the mutual aid relationships of hospitals in the network, we formulate the problem as a game theoretic model. We show that the response sets are nested and we provide an algorithm to obtain numerically the Nash Equilibrium solution of this game. We illustrate the structure of the model on a two-hospital example and perform sensitivity analysis with respect to parameters of our model.**

I. INTRODUCTION

With the recurrence of the Influenza A virus subtype H5N1 (also called *avian flu* or *bird flu*) in recent years, many public health experts believe that the world is closer to a pandemic event than at any time since the Hong Kong flu (H3N2) outbreak of 1968. A pandemic may happen quickly once the virus achieves efficient human to human transmission. Such an event is a threat to healthcare facilities and a challenge to public health officials across the country. An influenza pandemic will likely cause a large surge of patients over a short period of time and therefore potentially overwhelm the healthcare systems that are already operating close to capacity. As a consequence, public health and healthcare decision-makers have started to institute flu pandemic response plans including the establishment of flu pandemic response procedures, protocols, and cooperation between government agencies and healthcare providers. Some government sources, such as the World Health Organization (WHO) and the the Department of Health and Human Services (HHS), provide general checklists and recommendations on how to achieve preparedness goals [1]. However, there are scarce evidence-based resources that can guide healthcare facilities to establish and maintain sufficient surge capacity, including preparing enough medical supplies for a flu pandemic.

The research reported in this paper addresses the issue of hospital stockpiling in preparation for a flu pandemic. We assume that hospitals in a regional network have mutual aid agreements under which each hospital may borrow or lend supplies when in a medical emergency. Therefore, a game theoretic framework is applied to this problem to capture the effect of supply sharing between hospitals on stockpile quantities. Moreover, since there are many uncertainties associated with how a flu pandemic may start, how fast it may spread and how severe it may become, we take a stochastic approach. We assume that the overall hospital surge demand of the region is determined by the flu attack rate which follows a certain probability distribution. The patient demand at each hospital in the network is a fraction of the total demand that is proportional to its size. We also assume that each hospital in the network has a pre-determined response level to a possible flu pandemic. We will show that this assumption leads to ordered so-called nested *response sets*, which we define below. A game forms when each hospital makes the decision of its stockpile level in order to minimize the expected total cost incurred with consideration of possible borrowing and lending, given the stockpiling decisions of other hospitals in the network. To differentiate this model from our previously developed hospital stockpiling game [2], we call it a game with pre-determined response levels.

The organization of this paper is as follows: Section II describes the game model of stockpiling for a flu pandemic and explores the properties of the model. Section III presents some numerical examples, and sensitivity analysis. Section IV concludes this paper, and suggests future work.

II. MODEL DESCRIPTION

A. Model Setup

Consider a network, H, of n hospitals participating in a medical supply stockpiling game. The decision variable for each hospital $i = 1, \ldots, n$ is its stockpile level, s_i. As mentioned earlier, these hospitals have mutual aid agreement in place under which each member in the network may share supplies with one another in an emergency. Let α be the clinical flu attack rate that determines the severity of a pandemic, defined as the percentage of clinical influenza illness cases per population [3]. We assume that the probability distribution of α is known with density function $f(.)$ and cumulative distribution function $F(.)$. Note that $F(0) = 0$.

P. C. DeLaurentis, School of Industrial Engineering, Purdue University, West Lafayette, IN 47907 USA (765-496-1950), chenp@purdue.edu)

E. Adida, Department of Mechanical & Industrial Engineering, University of Illinois at Chicago, Chicago, IL 60607 USA, (elodie@uic.edu)

M. Lawley is with the Weldon School of Biomedical Engineering, Purdue University, West Lafayette, IN 47907 USA, (malawley@purdue.edu)

The set of pandemic scenarios, P, is defined as the collection of scenarios induced by all possible attack rates.

$$P = \{\alpha \mid \alpha \geq 0,\ g(\alpha) \leq Bound\}$$

where $Bound$ is an upper limit on the attack rate and $g(.)$ is an increasing bounding function.

Assumption 1: We assume that the overall surge demand of the hospital network, $D(.)$, is a continuous, strictly increasing function of α, and that each hospital i in the network receives a fraction of $D(\alpha)$ in proportion to its hospital size, z_i.

Thus, the surge demand for i under pandemic scenario of attack rate α can be expressed as $d_i(\alpha) = a_i \cdot D(\alpha)$, where $a_i = \frac{z_i}{\sum_{j=1}^{n} z_j}$. We observe that $D(.)$ is invertible and its inverse $D^{-1}(.)$ is continuous and strictly increasing.

Assumption 2: We assume that hospital i pre-determines a bound, $B_i > 0$, as the maximum surge demand for which it plans.

We call response set for hospital i, denoted R_i, the subset of pandemic scenarios such that the surge demand at i is within the pre-determined bound. By definition,

$$R_i = \{\alpha \mid d_i(\alpha) \leq B_i\} = \{\alpha \mid D(\alpha) \leq \frac{B_i}{a_i}\} = [0, D^{-1}(\frac{B_i}{a_i})].$$

The probability of a pandemic occurring within i's response set is

$$Pr(\alpha \in R_i) = F(D^{-1}(\frac{B_i}{a_i})).$$

B. Nested Response sets

The following result shows that the response sets are nested.

Lemma 1: (Nested Response sets) We have either $R_i \subset R_j$ or $R_j \subset R_i$.

Proof: We showed above that $R_i = [0, D^{-1}(\frac{B_i}{a_i})]$ and $R_j = [0, D^{-1}(\frac{B_j}{a_j})]$. If $\frac{B_i}{a_i} \leq \frac{B_j}{a_j}$, then since D^{-1} is strictly increasing $D^{-1}(\frac{B_i}{a_i}) \leq D^{-1}(\frac{B_j}{a_j})$ and $R_i \subset R_j$. Similarly, if $\frac{B_i}{a_i} \geq \frac{B_j}{a_j}$ then $R_j \subset R_i$. ∎

Without loss of generality, we assume that the hospitals are ordered such that

$$B_1/a_1 \leq \ldots \leq B_n/a_n.$$

Then clearly $R_1 \subset R_2 \subset \ldots \subset R_n$. Given a pandemic scenario α, we define the responding hospital set to be

$$R(\alpha) = \{i \in H \mid \alpha \in R_i\} = \left\{i \in H \mid \alpha \leq D^{-1}(B_i/a_i)\right\},$$

and the under-responding hospital set to be

$$\begin{aligned} N(\alpha) &= R(\alpha)^c = \{i \in H \mid \alpha \notin R_i\} \\ &= \left\{i \in H \mid \alpha > D^{-1}(B_i/a_i)\right\}. \end{aligned}$$

If $i \in R(\alpha)$, then $\alpha \in R_i$ and thus $\alpha \in R_j\ \forall j > i$.

We observe that

$$Pr(i \in R(\alpha)) = Pr(\alpha \in R_i) = F(D^{-1}(\frac{B_i}{a_i})).$$

Figure 1 is a representation of the nested response sets and the responding hospital set.

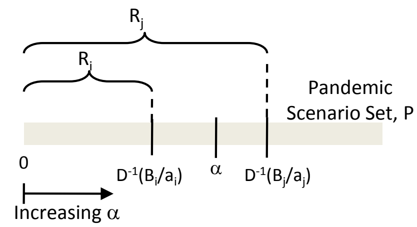

Fig. 1. Illustration of nested response sets. For i and j such that $D^{-1}(B_i/a_i) < \alpha < D^{-1}(B_j/a_j)$, we know that $j \in R(\alpha)$ and $i \in N(\alpha)$.

C. Demand redistribution

We assume that if a pandemic occurs outside of a hospital's response set, i.e. if the surge demand exceeds the pre-determined bound, the hospital's unmet demand (demand in excess of the bound) is redistributed among hospitals in the network whose response set contains the pandemic realization. However, if the severity of a pandemic lies within a hospital's response set, then there is no demand redistribution from this hospital. Let the demand redistributed from $j \in H$ be $r_j(\alpha)$. The two cases described above can be mathematically written as follows.

- If $j \in N(\alpha)$, then $d_j(\alpha) > B_j$ and $r_j(\alpha) = d_j(\alpha) - B_j$.
- If $j \in R(\alpha)$, then $d_j(\alpha) \leq B_j$ and $r_j(\alpha) = 0$.

D. Receipt of redistributed demand

As described above, the hospitals in the responding set receive redistributed demand from under-responding hospitals. For $i \in R(\alpha)$, the proportion of the total redistributed demand that i receives is

$$w_i(\alpha) = \frac{z_i}{\sum_{j \in R(\alpha)} z_j},$$

where we recall that z_i is the size of hospital i. Clearly, $w_k(\alpha) = 0,\ \forall k \in N(\alpha)$. Therefore the redistributed demand that i receives is

$$\begin{aligned} x_i(\alpha) &= w_i(\alpha) \cdot \sum_{j \in N(\alpha)} r_j(\alpha) \\ &= w_i(\alpha) \cdot \sum_{j \in N(\alpha)} (d_j(\alpha) - B_j) \end{aligned}$$

E. Inventory surplus & deficit

Hospital i's inventory surplus ($y_i^+(s_i)$) and deficit ($y_i^-(s_i)$) are defined as the following:

- If hospital i is responding, i.e., $i \in R(\alpha)$, then

$$\begin{aligned} y_i^-(s_i) &= \max\{0,\ d_i(\alpha) + x_i(\alpha) - s_i\} \\ y_i^+(s_i) &= \max\{0,\ s_i - d_i(\alpha) - x_i(\alpha)\} \end{aligned}$$

- If hospital i is under-responding, i.e., $i \in N(\alpha)$.

$$\begin{aligned} y_i^+(s_i) &= \max\{0,\ s_i - B_i\} \\ y_i^-(s_i) &= \max\{0,\ B_i - s_i\} \end{aligned}$$

Note that we make the dependency on s_i explicit, but not the dependency on α for ease of exposition.

We further define the system-wide inventory surplus and deficit as

$$Y^+(s) = \sum_{i \in H} y_i^+(s_i)$$
$$Y^-(s) = \sum_{i \in H} y_i^-(s_i)$$

We observe that the system-wide inventory surplus and deficit depend on the stockpiling decisions of all hospitals (as well as, implicitly, on α). Clearly, if $Y^-(s) > Y^+(s)$ then there is a deficit overall (i.e. the overall stockpile is lower than the overall demand), while if $Y^-(s) < Y^+(s)$ then there is a surplus overall. We assume that supplies are distributed evenly across all hospitals. In other words, if there is a deficit overall, hospital i with an inventory surplus will loan all of its surplus $y_i^+(s_i)$, and hospital j with an inventory deficit will receive $y_j^-(s_j)\frac{Y^+(s)}{Y^-(s)}$. If there is a surplus overall, hospital i with an inventory surplus will loan $y_i^+(s_i)\frac{Y^-(s)}{Y^+(s)}$, and hospital j with an inventory deficit will receive $y_j^-(s_j)$.

We observe that

$$Y^-(s) > Y^+(s) \;\Leftrightarrow\; \sum_{i=1}^{n} d_i(\alpha) > \sum_{i=1}^{n} s_i \;\Leftrightarrow\; D(\alpha) > \sum_{i=1}^{n} s_i$$

and therefore

$$Pr(Y^-(s) > Y^+(s)) = 1 - F(D^{-1}(\sum_{i=1}^{n} s_i)).$$

F. Objective function

The objective of each hospital is to minimize the expected total cost incurred by stockpiling medical supplies. The total cost includes the cost of purchasing the supplies, the holding cost until the start of the pandemic and the cost of borrowing supplies from other hospitals in the network if necessary. A hospital that lends supplies receives a compensation. Moreover, we assume that hospitals get penalized for not having enough supplies to meet the demand, and for setting a pre-set limit too low. Hospital i's total stockpile cost is then given by

$$\begin{aligned} J_i(s_i, s_{-i}) = [&\text{Stockpile Purchase \& Holding Cost} \\ &- \text{Supply Sharing Income} \\ &+ \text{Supply Borrowing Cost} \\ &+ \text{Supply Shortage Penalty} \\ &+ \text{Under Prepared Penalty}] \end{aligned} \tag{1}$$

where s_{-i} denotes the stockpile levels of all hospitals except i. The optimization problem faced by hospital i depends on the decisions made by all other hospitals, hence it is called a best-response problem. Our goal is to find a Nash equilibrium for this game, i.e. a set of stockpile levels for all hospitals such that they are all at their best response and have no incentive to deviate.

The best response problem of hospital i can be written:

$$\begin{aligned} &\min_{s_i} E[J_i(s_i, s_{-i})] \\ =&\min_{s_i} E\Big[cs_i + hTs_i - b_1 \min\Big\{y_i^+(s_i),\ y_i^+(s_i)\frac{Y^-(s)}{Y^+(s)}\Big\} \\ &\qquad + b_2 \min\Big\{y_i^-(s_i),\ y_i^-(s_i)\frac{Y^+(s)}{Y^-(s)}\Big\} \\ &\qquad + p_s \max\Big\{0,\ y_i^-(s_i) - y_i^-(s_i)\frac{Y^+(s)}{Y^-(s)}\Big\} \\ &\qquad + p_r \max\{0, d_i(\alpha) - B_i\}\Big] \\ =&\min_{s_i} (c + hE[T])s_i - b_1 E\Big[\min\Big\{y_i^+(s_i),\ y_i^+(s_i)\frac{Y^-(s)}{Y^+(s)}\Big\}\Big] \\ &\quad + b_2 E\Big[\min\Big\{y_i^-(s_i),\ y_i^-(s_i)\frac{Y^+(s)}{Y^-(s)}\Big\}\Big] \\ &\quad + p_s E\Big[\max\Big\{0,\ y_i^-(s_i) - \Big(y_i^-(s_i)\frac{Y^+(s)}{Y^-(s)}\Big)\Big\}\Big] \\ &\quad + p_r E[\max\{0, d_i(\alpha) - B_i\}] \end{aligned} \tag{2}$$

The symbols used in this model are listed in Table I. Expression (2) is too complex to solve in closed-form because of the terms such as $E\Big[\min\Big\{y_i^+(s_i),\ y_i^+(s_i)\frac{Y^-(s)}{Y^+(s)}\Big\}\Big]$ involve intricate conditional expectations.

TABLE I
NOTATION USED IN THIS MODEL

Symbol	Explanation
H	Set of hospitals in the network of interest
n	number of hospitals in H
s_i	Stockpile level of hospital i; decision variable
α	Flu gross attack rate (random variable)
$F(.)$	cumulative distribution function for α
$f(.)$	probability density function for α
P	All possible pandemic scenarios
$D(.)$	Overall surge demand of the hospital network
z_i	Size of hospital i
$d_i(.)$	Demand of hospital i
a_i	Fraction of overall demand at hospital i
B_i	Hospital i's pre-set response level
R_i	Response set of hospital i
$R(.)$	Set of responding hospitals
$N(.)$	Set of under-responding hospitals
$r_i(.)$	Amount to be redistributed to others
$w_i(.)$	Proportion of demand redistributed to i
$x_i(.)$	Redistributed demand that i receives
$y_i^+(.)$	Hospital i's inventory surplus
$y_i^-(.)$	Hospital i's inventory deficit
$Y^+(.)$	System-wide inventory surplus
$Y^-(.)$	System-wide inventory deficit
$J_i(.)$	Total stockpiling cost of hospital i
b_1	Unit cost charged by lending hospitals
b_2	Unit cost charged to borrowing hospitals
c	Unit purchasing cost
h	Unit holding cost
p_s	Unit penalty cost due to supply shortage
p_r	Unit penalty cost due to under preparedness
T	Time until the onset of next flu pandemic (random variable)

G. Solution method

1) Estimating objective function: Since no closed-form expression of the best-response objective function can be obtained, we adopt a simulation approach to solve the problem numerically. The primary uncertainty in the stockpiling game is the severity of a flu pandemic which is characterized by its clinical attack rate, α. As mentioned earlier, we assume α follows a certain distribution. Thus, the simulation method starts by sampling α according to its distribution. The steps to estimate the expected total stockpiling cost for a given vector $s = (s_1, \ldots, s_n)$ are shown in Table II.

TABLE II
STEPS TO ESTIMATE EXPECTED TOTAL STOCKPILING COST

1. Set sample size to M.
2. For k=1 to M
   Sample $\alpha = \alpha^k$ according to its distribution.
   Calculate $J_i^k(s)$ as the value of $J_i(s)$ for $\alpha = \alpha^k$.
   End For
3. Calculate expected $J_i(s)$ by averaging all $J_i^k(s)$,
   i.e. $E[J_i(s)] = \sum_{k=1}^{M} J_i^k(s)$

For a given s_{-i}, once the expected objective function, $E[J_i(s_i, s_{-i})]$, is obtained for a large number of values of s_i, the best response can be found as the stockpile level, $s_i^*(s_{-i})$, that minimizes $E[J_i(s_i, s_{-i})]$ with respect to s_i.

2) Finding a Nash equalibrium: A Nash equilibrium is a commonly known concept for defining the outcome of a noncooperative game. it is defined as a profile of strategies at which each player's strategy is an optimal response to other players' strategies [4]. Mathematically, it is the fixed point of the best response mapping. A unique Nash equilibrium is especially useful since it helps predict the outcome of the game. However, not all games have a Nash solution, and some games may have multiple Nash equilibria.

For a two-hospital game as the example shown in the next section, a Nash equilibrium is simply the intersection of the best response functions of the two hospitals involved. If the Nash equilibrium cannot be seen on a plot, a solution can be obtained by some algorithms designed to search for a Nash equilibrium. One example is the *iterative algorithm*. In such an algorithm, players update their strategies iteratively by selecting unilaterally their best responses in turn until adjustments are no longer needed and an equilibrium is thus reached. At each iteration, each one of the players sequentially selects his or her best strategy responding to other players' decisions in the previous stage. This iterative process stops when each one of the players' strategies is within a very small tolerance of the best response, and the current set of strategies is deemed the equilibrium solution [4], [5].

III. NUMERICAL EXAMPLES

In order to illustrate the mathematical properties and concepts of this model, a two-hospital stockpiling game is set up and shown in this section. We let the clinical attack rate, α, follows a triangular distribution with (lower limit, upper limit, mode) = (0, 35, 7) since we assume that α is likely to range from 0% to 35% (as in the great pandemic in 1918) with most probable case of a seasonal-flu-like scenario which has α=7%. The values of hospital characteristics and stockpile costs used throughout this section are shown in Table III.

TABLE III
DATA USED IN THE 2-HOSPITAL GAME EXAMPLE

c	5
h	1
b_1	6
b_2	7
p_s	60
p_r	60
$B = [B_1 B_2]$	[150 50]
$E[T]$	12
$z = [z_1 z_2]$	[120 30]
Baseline demand @ $\alpha = 7\%$	200

A. Illustration of objective function

The first illustration is the objective function. We consider a game with two hospitals indexed $i = 1, 2$. Figure 2 shows that the objective function of hospital 1, $J_1(.)$, varies with the value of the stockpile level of hospital 2, s_2. Note that $J_1(.)$ appears to be convex in s_1. For a given s_2, the best response $s_1^*(s_2)$ of hospital 1 is the value of s_1 at which $J_1(s_1, s_2)$ is minimized with respect to s_1. We notice that as s_2 increases, the overall value of $J_1(.)$ decreases and $s_1^*(s_2)$ decreases. This is reasonable since the more hospital 2 stockpiles, the more hospital 1 can expect to borrow from hospital 2 when needed, which reduces hospital 1's overall cost and optimal stockpile quantity.

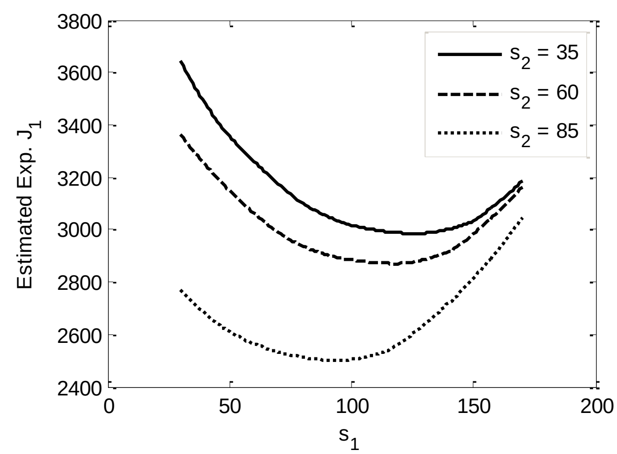

Fig. 2. Example of the estimated expected objective function of hospital 1, $E[J_1(s_1, s_2)]$, as a function of its stockpile level s_1 and for 3 possible values of hospital 2's stockpile level, s_2, in a 2-hospital game.

B. Illustration of best response function

As mentioned earlier, a Nash equilibrium corresponds to the set pf decisions such that all players are at their best response. Geometrically, a Nash equilibrium lies is the intersection of the best response functions. In Figure 3, the solid line is the best response of hospital 1, $s_1^*(s_2)$, as a function of the stockpile of hospital 2, while the dashed line is the best response of hospital 2, $s_2^*(s_1)$, as a function of hospital 1's stockpile level. It clearly shows that there is exactly one intersection of these two functions. A closer look at the plot reveals that the two best response functions intersect at $(s_1, s_2) = (150, 50)$ which is a Nash equilibrium in this example.

Note that the solution is the stockpile level of one medical item under the assumed cumulative surge demand predicted throughout the course of a flu pandemic. These numbers can easily be scaled up or down for different types of medical supplies which may have different replacement rates depending on common practice or the standards of medical care.

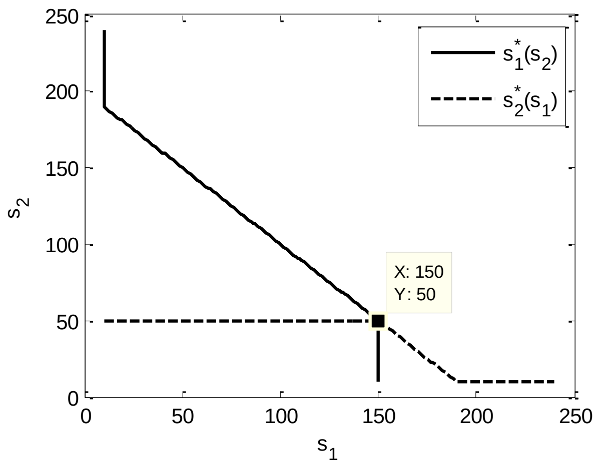

Fig. 3. Example of best response functions in a 2-hospital game. $s_1^*(s_2)$ is the best response of hospital 1 as a function of hospital 2's stockpile level. Similarly, $s_2^*(s_1)$ is the best response of hospital 2 as a function of hospital 1's stockpile level.

C. Sensitivity analysis

We recognize that it is generally difficult to estimate some of the parameters used in this game model such as the expected time to the onset of the next flu pandemic, the penalty cost of shortage and the penalty cost of under-preparedness. Therefore, it is useful to perform a sensitivity analysis of the best response solutions with respect to these values. It will help us understand the behavior of the hospital stockpiling game model. We select two most influencing factors in the model to present in this paper.

1) Expected time to onset of a flu pandemic: As mentioned before, there are many uncertainties associated with a possible pandemic outbreak. Since the onset of such a disease depends greatly on the characteristics of the virus, scientists and/or public health experts cannot predict precisely when it will occur. Though unpredictable, the expected time to the onset of the next pandemic, $E[T]$, is an important parameter which affects the inventory holding cost incurred for a hospital supply stockpile. Figure 4 shows that the Nash solution varies with different values of $E[T]$.

It is clear that when the expected time to the next pandemic is longer, hospitals will tend to stockpile less. This is so because if hospitals expect to store supplies in the warehouse for a longer period of time (thus spend more money in storage), they likely will stockpile less in order to minimize the storage (holding) cost. The relationship between $E[T]$ and the stockpile level is also valuable information for public health policy makers. It proves that when communicating with healthcare providers or even the general public on pandemic planning matters, an urgency in the message may be key to timely and sufficient response to the task.

2) Penalty cost of supply shortage: We also recognize the difficulty of assigning a value to the penalty cost due to supply shortage, p_s, in this model. Figure 5 shows the relationship between p_s and hospital stockpile levels indicating that when a higher penalty cost is assessed (such as when a hospital is more risk averse), hospitals will maintain a larger quantity of inventory. When the consequence of not meeting surge demand is low, i.e., p_s is small, hospitals will stock a lesser amount.

IV. CONCLUSIONS AND FUTURE WORK

A. Conclusions

In this research, we adopt a game theoretical approach to study the stockpiling decisions of a network of hospitals in anticipation of a flu pandemic. A hospital's decision on the stockpile level of medical supplies depends not only on the uncertainties associated with the possible next flu pandemic, but also on its mutual aid relationship with other partner hospitals in the network. Since this hospital stockpiling game model does not have a closed-form expression for each hospital's best response, we use a simulation approach to estimate the objective functions and best responses by sampling flu attack rate, α. We illustrate on a two-hospital example the Nash equilibrium solution. We also show some sensitivity analysis results indicating hospital stockpile levels vary with different expected time to the next pandemic and penalty cost due to supply shortage.

B. Future work

We envision the future work of this model will include the presentation of a n-hospital game (where $n > 2$) and its Nash solutions. The sensitivity of the solutions should also be tested and compared with those in a 2-hospital game.

REFERENCES

[1] The Department of Health and Human Services, PandemicFlu.gov. (URL: http://pandemicflu.gov/index.html)

[2] P. C. DeLaurentis, E. Adida and M. Lawley. A Game Theoretical Approach for Hospital Stockpile in Preparation for Pandemics. *Proceedings of the 2008 Industrial Engineering Research Conference, J. Fowler and S. Mason, eds*. Vancouver, Canada. May 2008.

[3] Martin I. Meltzer, Nancy J. Cox and Keiji Fukuda, "The economic impact of pandemic influenza in the United States: priorities for intervention", *Emerging Infectious Diseases*, vol. 5, no. 5, pp. 659–671, 1999.

[4] D. Fudenberg and J. Tirole. 1991. *Game Theory*. The MIT Press, Cambridge, MA.

[5] Daniel M. Reeves and Michael P. Wellman. Computing Best-Response Strategies in Infinite Games of Incomplete Information. In *Twentieth Conference on Uncertainty in Artificial Intelligence*, pp. 470-478, 2004.

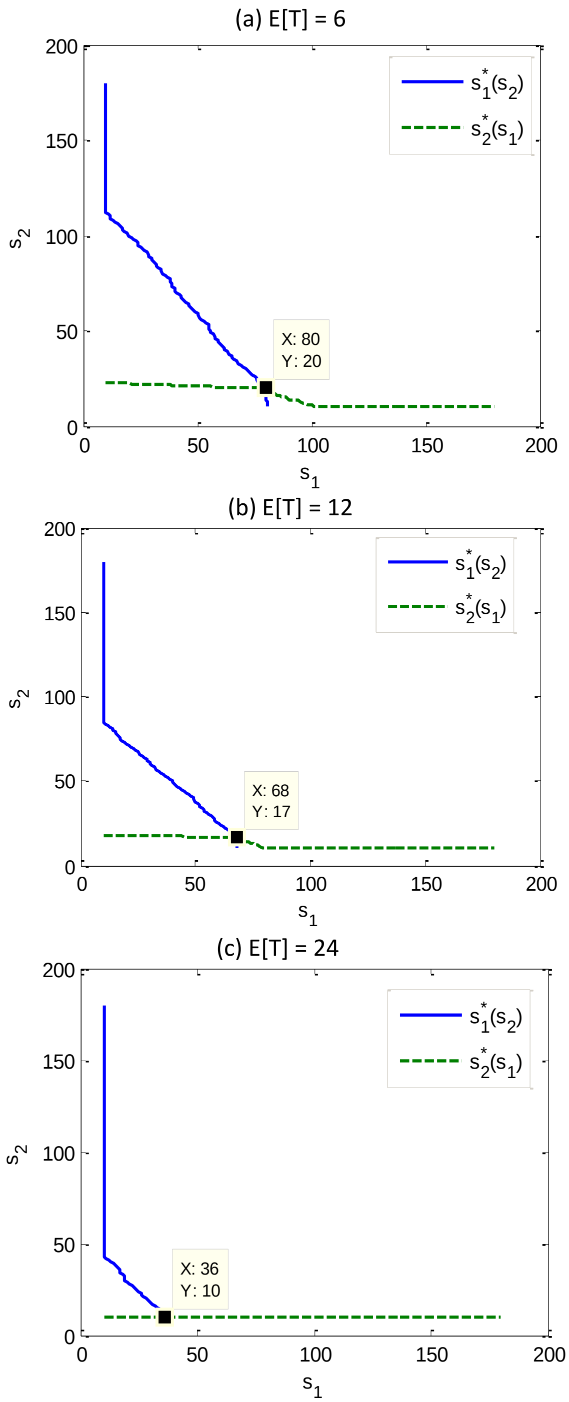

Fig. 4. The solution of a 2-hospital game varies with the expected time to the onset of next pandemic, $E[T]$. As $E[T]$ increases, hospital stockpile levels decrease.

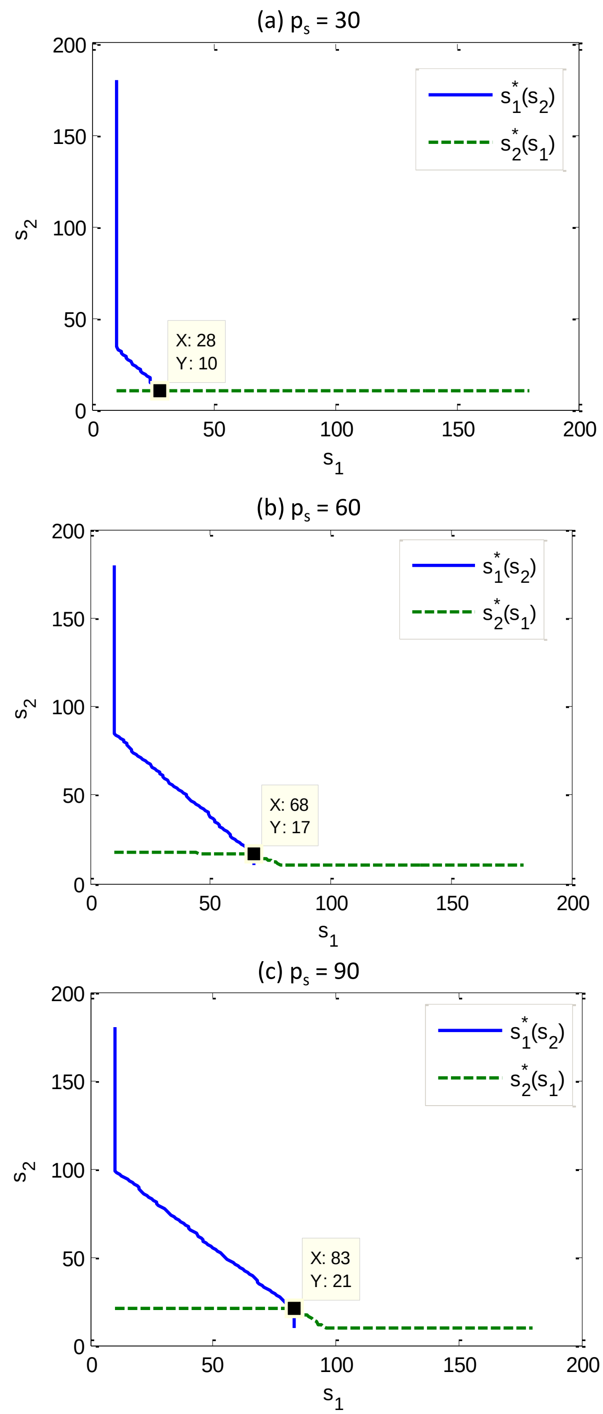

Fig. 5. The solution of a 2-hospital game varies with the penalty cost for shortage, p_s. As p_s increases, hospitals stockpile more supplies.