

Hospital Stockpiling Problems with Inventory Sharing

Eric Lofgren[†] and Anil Vullikanti^{*,†}

[†]Biocomplexity Institute of Virginia Tech, Virginia Tech

^{*}Department of Computer Science, Virginia Tech
Email: {lofgrene, akumar}@vbi.vt.edu

Abstract

Hospitals are typically optimized to operate near capacity, and there are serious concerns that our healthcare system is not prepared for the next pandemic. Stockpiles of different supplies, e.g., personal protective equipments (PPE) and medical equipment, need to be maintained in order to be able to respond to any future pandemics. Large outbreaks occur with a low probability, and such stockpiles require big investments. Further, hospitals often have mutual sharing agreements, which makes the problem of stockpiling decisions a natural game-theoretical problem. In this paper, we formalize hospital stockpiling as a game-theoretical problem, HSTOCKPILE. We use the notion of pairwise Nash stability as a solution concept for this problem, and characterize its structure. We show that stable strategies can lead to high unsatisfied demands in some scenarios, and stockpiles might not be maintained at all nodes. We show that stable strategies and the social optimum can be computed efficiently.

1 Introduction

Large disease outbreaks can place a very significant burden on the healthcare infrastructure, as recent Ebola and 2009 H1N1 influenza epidemics have shown. Hospitals and other healthcare facilities are optimized for a baseline level of patient demand. Even a slight increase in the infection rate can lead to a significant surge in patient visits, e.g., as in the case of the first 40 weeks of the 2009 Flu outbreak, when the number of patient visits increased by 4% over the baseline (DeLaurentis, Adida, and Lawley 2011). Handling such a surge would require significant level of planning for disaster events and investments to handle the surge in patient demand—this includes medical equipment (e.g., personal protective equipments (PPE), patient care supplies), medications such as anti-virals, hospital facilities and healthcare personnel (ton). The CDC and other regional and national public health agencies have developed guidelines for investments in stockpile levels. However, it is challenging for healthcare companies to make huge investments for handling such a surge, especially because of its variable and stochastic nature. In order to amortize the cost of such investments, and since the patient demand at an individual hospital might vary widely, it is common for hospitals to engage in sharing

agreements (e.g., (Stier and Goodman 2007)). This allows them to buy or sell supplies, based on mutual agreements. A fundamental problem is how should hospitals decide on their initial stockpile levels, and how should they trade supplies based on an outbreak? How should the government design incentives (or penalties), so that the health infrastructure is best prepared for pandemic outbreaks?

These are natural game-theoretical problems and DeLaurentis et al. (DeLaurentis, Adida, and Lawley 2011) develop the first non-cooperative game formulation for the hospital stockpiling problem. In their formulation, each hospital i needs to decide on its stockpile level s_i . The epidemic outbreak is assumed to be a distribution over a small number of scenarios, with $D(i, \ell)$ denoting the patient demand at hospital i in scenario ℓ . A hospital i has surplus of $D(i, \ell) - s_i$ if $s_i > D(i, \ell)$ and a deficit of $s_i - D(i, \ell)$ if $s_i < D(i, \ell)$. It is assumed that surplus at any hospital i can be sent to any other hospital j , at some cost. Further, a hospital i has to pay a penalty p for each unit of patient demand that is not satisfied. (DeLaurentis, Adida, and Lawley 2011) study the structure of Nash equilibria in such a game, and find that they can be fairly inefficient.

While the work of (DeLaurentis, Adida, and Lawley 2011) provides the first formal approach to study the issues of stockpiling, they abstract out many realistic aspects—one of these is the network structure among the hospitals, so that their might be a limit $\text{cap}(i, j)$ on how much supplies can be traded between hospitals i and j . This might also affect the cost C_{ij} of trading supplies between the hospitals. As studied in (Lee et al. 2011; Donker, Wallinga, and Grundmann 2010; Simmering et al. 2015), there is a clear network structure among hospitals, which plays a role in the patient referrals across hospitals. The role of this network structure has been studied extensively in studying and controlling the spread of hospital acquired infections. Such interactions require prior agreements, which will play a role in sharing stockpiles, in the event of pandemic outbreaks.

In this paper, we extend the work of (DeLaurentis, Adida, and Lawley 2011) to incorporate a much more realistic stockpiling problem with network based sharing constraints. This takes the first step towards addressing one of the extensions suggested by them on considering “... implementable contracts among hospitals and possibly involving government in a transfer payment scheme ...”. Our contributions are:

1. We formalize the hospital stockpiling game problem (HSTOCKPILE) with network capacity constraints as a non-cooperative game. This captures much more realistic network constraints.
2. We use the notion of pairwise Nash stability and characterize stable strategies. We find that generally deficits remain in stable strategies, depending on the stockpiling costs and penalty for deficits. The nodes that select to stockpile form a dominating set in the graph, under certain conditions. Further, link capacities have a very significant impact on the efficiency of stable solutions.
3. We show that any local optimum is a stable solution, and that a social optimum (a minimum cost stable solution) can be found using a linear programming approach. We use this method to find the solutions in a hospital network in the state of North Carolina, and observe that there are significant levels of deficits in practice. Our results suggest the need for improved strategies, such as subsidies to help reduce the deficits and ensure better preparedness.

2 Preliminaries and Model

Let V denote a set of hospitals in a region, which needs to make decisions about their individual stockpile level. We refer to each hospital $i \in V$ as a node. Since epidemic outbreak is a stochastic process, as in (DeLaurentis, Adida, and Lawley 2011), we assume a set \mathcal{D} of scenarios, with a patient demand of $D(i, \ell)$ for hospital i in scenario $\ell \in \mathcal{D}$. Let q_ℓ denote the probability of scenario ℓ . Let C_i denote the per unit cost of stockpiling for hospital i , and C_{ij} the per unit cost of sharing between the hospitals i and j . We assume $C_{ij} = C_{ji}$. \mathbf{C} denotes the vector of all costs. We let $\text{cap}(i, j)$ denote the maximum amount of supplies that can be sent from i to j . We let $G = (V, E)$ denote the hospital network with $E = \{(i, j) : \text{cap}(i, j) > 0\}$ denoting the set of edges with positive trading capacity.

The main objective is to determine the stockpile level $s(i)$ for each hospital $i \in V$. Let C_i denote hospital i 's per unit cost of maintaining the stockpile. Since the patient demand $D(i, \ell)$ might vary with the scenario ℓ , there might be scenarios ℓ for which $D(i, \ell) > s(i)$. In this case, hospital i can buy supplies from other hospitals j , at a per-unit cost of C_{ji} . Let $s(j, i, \ell)$ and $s(i, j, \ell)$ denote the amount bought by hospital i from hospital j , or sold by hospital i to hospital j , for scenario ℓ , respectively. These must satisfy the capacity constraints, so that $s(i, j, \ell) \leq \text{cap}(i, j)$ for all i, j, ℓ . Let $s(i, \ell) = s(i) - \sum_{j \neq i} s(i, j, \ell)$ denote the amount of supplies available at node i in scenario ℓ . Also, let $s_{\text{out}}(i, \ell) = \sum_{j \neq i} s(i, j, \ell)$ and $s_{\text{in}}(i, \ell) = \sum_{j \neq i} s(j, i, \ell)$ denote the total amount of supplies sold and purchased by i , respectively. Then,

$$\text{def}(i, \ell) = \max\{D(i, \ell) - s(i, \ell) + s_{\text{in}}(i, \ell), 0\}$$

is the patient demand that is not satisfied at node i in scenario ℓ (referred to as the deficit at node i in scenario ℓ). We assume there is a penalty p for each unit of demand that is not satisfied. We assume that for each i, j ,

$$C_i \leq C_{ij} \leq p,$$

so that node i is always better off purchasing supplies in case of a deficit, instead of paying a penalty.

The HSTOCKPILE game. An instance of HSTOCKPILE consists of a tuple $(G, \text{cap}, \mathbf{C}, \mathbf{D}, p)$. We consider a non-cooperative setting, in which the strategy of each hospital is the initial stockpile $s(i)$, and the amount $s(j, i, \ell)$ or $s(i, j, \ell)$ it buys or sells from other hospitals j . The game is played in the following manner:

1. Before the start of the epidemic: initially each node i decides its stockpile level $s(i)$. Each pair of nodes i, j decide on the amounts $s(i, j, \ell)$ and $s(j, i, \ell)$, for each scenario ℓ .
2. The epidemic spreads as a stochastic process, and each scenario ℓ occurs with probability q_ℓ , leading to a demand $D(i, \ell)$ for hospital i . If scenario ℓ occurs, nodes i, j trade $s(i, j, \ell)$ or $s(j, i, \ell)$ amount between themselves.
3. Node i incurs cost

$$\begin{aligned} \text{cost}(i, s) = & s(i)C_i + \sum_{\ell} q_\ell \sum_{j \neq i} s(j, i, \ell)C_{ji} \\ & - \sum_{\ell} q_\ell \sum_{j \neq i} s(i, j, \ell)C_{ij} + \sum_{\ell} q_\ell p \max\{0, \text{def}(i, \ell)\} \end{aligned}$$

For succinctness, we use $s(\cdot)$ to denote the entire strategy profile, consisting of the stockpile amounts $s(i)$ and the traded amounts $s(i, j, \ell)$ in each scenario.


Stable strategies and social optimum

For a strategy profile $s(\cdot)$, the pairwise strategy $s(i, j, \ell)$ needs to be jointly decided by both i and j . Therefore, the standard notion of Nash equilibria (NE), which is based on no single individual having incentive to deviate unilaterally (Leyton-Brown and Shoham 2008) is not suitable here. Instead, we use a notion of *pairwise Nash stability*, which has been studied for network formation games (see, e.g., (Jackson 2008; Blume et al. 2012))—in this notion, a strategy is stable if no pair of players have incentive to deviate. We apply this notion to the HSTOCKPILE problem. We say that a strategy $s(\cdot)$ is *stable* if the following conditions hold:

1. No node i can reduce its cost by unilaterally changing its stockpile $s(i)$, while keeping all the $s(i, j, \ell)$ unchanged, for $j \neq i$. Changing $s(i)$ without changing $s(i, j, \ell)$ or $s(j, i, \ell)$ has the effect of changing $s(i, \ell)$.
2. There is no pair of nodes i, j such that changing $s(i, j, \ell)$, $s(j, i, \ell)$, $s(i)$ or $s(j)$, while all other components of $s(\cdot)$ are fixed, does not increase the cost for either i or j , and reduces the cost for at least one of them.

Formally, strategy $s(\cdot)$ is stable if:

1. For all $i \in V$, we have $\text{cost}(i, s') \geq \text{cost}(i, s)$, where $s'(\cdot)$ is any strategy such that
 - (a) For all $j \neq i$: $s'(j) = s(j)$
 - (b) For all $i' \neq j'$, $s'(i', j', \ell) = s(i', j', \ell)$.
2. For all $i, j \in V$, either (a) $\text{cost}(i, s') > \text{cost}(i, s)$ or $\text{cost}(j, s') > \text{cost}(j, s)$, or (b) $\text{cost}(i, s') \geq \text{cost}(i, s)$ and $\text{cost}(j, s') \geq \text{cost}(j, s)$, where $s'(\cdot)$ is any strategy such that



ℓ	$D(1, \ell)$	$D(2, \ell)$
1	0	200
2	100	300
3	200	0
4	300	100

Figure 1: Example of an instance of HSTOCKPILE: $V = \{1, 2\}$, $E = \{(1, 2)\}$. D consists of four scenarios, as shown in the table, each with probability $1/4$. The table gives the patient demands $D(i, \ell)$ are shown for each $i = 1, 2$. We have $C_i = 1$ for $i = 1, 2$, $C_{12} = 1.5$ and $p = 2$.

- (a) For all $i' \neq i, j$, $s'(i') = s(i')$
- (b) For all $\{i', j'\} \neq \{i, j\}$, $i' \neq j'$: $s'(i', j', \ell) = s(i', j', \ell)$.

We define the *social optimum* as a strategy $s(\cdot)$ that has the minimum cost, over the space of all possible strategies—this might, in general, not be a stable strategy. Therefore, the cost of a stable strategy relative to the cost of the social optimum is an important measure, and the maximum such ratio over all possible stable strategies is known as the *price of anarchy* (Koutsoupias and Papadimitriou 1999).

Example. Figure 1 gives a simple example of HSTOCKPILE. First, consider the case $\text{cap}(1, 2) = 0$, which corresponds to the two nodes being isolated. Consider a strategy $s(\cdot)$ with $s(1) = s(2) = 100$ and $s(i, j, \ell) = 0$. We have $\text{def}(1, 1) = \text{def}(1, 2) = 0$, $\text{def}(1, 3) = 100$, $\text{def}(1, 4) = 200$. Therefore, in this case, $\text{cost}(1, s) = 1 \cdot 100 + \frac{2}{4}(0 + 0 + 100 + 200) = 250$. Similarly, $\text{cost}(2, s) = 250$. For strategy $s'(\cdot)$ with $s'(1) = s'(2) = 50$, and $s'(i, j, \ell) = 0$ for all i, j, ℓ , we have $\text{def}(1, 1) = 0$, $\text{def}(1, 2) = 50$, $\text{def}(1, 3) = 150$, $\text{def}(1, 4) = 250$. Therefore, $\text{cost}(1, s') = 50 + \frac{2}{4}(0 + 50 + 150 + 250) = 275$. Similarly, $\text{cost}(2, s') = 275$. It can be verified that $s(\cdot)$ is the social optimum in this case. Next, consider the case where $\text{cap}(1, 2) = 50$. Let $s''(\cdot)$ be a strategy defined in the following manner: $s''(1) = s''(2) = 100$, $s''(1, 2, 1) = 50$, $s''(2, 1, 3) = 50$. In this case, we have $\text{cost}(1, s'') = 1 \cdot 100 + \frac{1}{4} \cdot 50 \cdot 1.5 + \frac{2}{4}(0 + 0 + 50 + 200) = 243.75$.

3 Related Work

Because of the major challenges posed by infectious diseases, all aspects of real time epidemiology and public health policy planning are very active areas of research (see, e.g., (Marathe and Vullikanti 2013)). We will focus our discussion here on different approaches for controlling the spread of epidemics, and especially those that involve game-theoretic and optimization based approaches.

Much of the efforts on controlling epidemics has focused on individual level interventions, e.g., distributing vaccinations and anti-virals (referred to as Pharmaceutical interventions (PI)), or methods to reduce transmission, such as by closing schools and social distancing (referred to as Non Pharmaceutical Interventions (NPI))—see (Meyers and Dimitrov 2010; Halloran et al. 2008) for discussion of these methods. Since individuals incur a cost in implementing such interventions, and they can be protected if enough other individuals follow them, these are naturally amenable to game

theoretical analysis. Much of this work has been based on differential equation methods, e.g., (Bauch and Earn 2004; Reluga and Galvani 2011). While these works enable rigorous analysis, they do not capture the realistic mixing patterns in social contact networks. This has motivated the study of vaccination games on network models, e.g., (Aspnes, Chang, and Yampolskiy 2006; Kumar et al. 2010; Saha, Adiga, and Vullikanti 2014). Computing efficient equilibria and socially optimal strategies turns out to be a very challenging problem, in general. An alternative approach has been to use spectral properties of networks for characterizing and controlling epidemic spread, e.g., (Saha, Adiga, and Vullikanti 2014; Omic, Orda, and Miegheem 2009; Trajanovski et al. 2015).

The focus of our paper is not on the decision making at an individual level, but at the level of hospitals and healthcare facilities. Hospital networks have been studied quite extensively, especially in the context of controlling the spread of Methicillin-Resistant Staphylococcus Aureus (MRSA) and other hospital-acquired infections, e.g., (Lee et al. 2011; Donker, Wallinga, and Grundmann 2010; Simmering et al. 2015). There also has been work on designing interventions for controlling epidemics at the level of hospital networks, e.g., (Prakash et al. 2013).

However, the problems of pandemic preparedness, such as stockpiling, remain relatively unexplored, especially from a formal game-theoretical perspective. The stockpiling problem is a special case of inventory modeling problems in economics and Operations Research, e.g., (Meca et al. 2004). To the best of our knowledge, the first formal analysis of these problems is by (DeLaurentis 2009; DeLaurentis, Adida, and Lawley 2011), who identifies many key problems in this topic. However, (DeLaurentis, Adida, and Lawley 2011) do not consider network constraints and only consider individual level decisions in their game formulations, which does not consider the incentives for a pair of hospitals to share supplies. Our paper is the first formalization of this problem.

4 Characterizing and computing stable strategies and the effect of network structure

We now discuss some structural properties of stable strategies, and the effects of network structure.

Lemma 1. Consider an instance $(G, \text{cap}, \mathbf{C}, \mathbf{D}, p)$ of HSTOCKPILE. For any stable strategy $s(\cdot)$, the following conditions hold:

1. For each i , we have $C_i = p \Pr[\text{def}(i, \ell) > 0] = p \sum_{\ell: \text{def}(i, \ell) > 0} q_\ell$.
2. For each i, j, ℓ such that $s(i, j, \ell) > 0$, then: (a) either $s(i, \ell) \leq D(i, \ell) - s_{in}(i, \ell)$ or $s(i, j, \ell) = \text{cap}(i, j)$ or $\text{def}(j, \ell) = 0$, and (b) either $\text{def}(i, \ell) = 0$ or $C_j \geq p \Pr[\text{def}(j, \ell) > 0]$.

Proof. The proof follows from the definition of stability. Recall that $\text{cost}(i, s) = C_i s(i) + \sum_{\ell} q_\ell p \max\{0, \text{def}(i, \ell)\} + z(i, s)$, where $z(i, s) = \sum_{\ell} q_\ell \sum_{j \neq i} s(j, i, \ell) C_{ji} -$




Figure 2: An instance of HSTOCKPILE with multiple stable solutions: $V = \{1, 2, 3\}$, $E = \{(1, 2), (2, 3)\}$. \mathcal{D} consists of two scenarios, as shown in the table, each with probability $1/2$. The table gives the patient demands $D(i, \ell)$ are shown for each $i = 1, 2$. We have $C_i = 1$ for $i = 1, 2, 3$, $C_{12} = C_{23} = 2$ and $p = 3$.

$\sum_{\ell} q_{\ell} \sum_{j \neq i} s(i, j, \ell) C_{ij}$. Let $s'(\cdot)$ be a strategy in which $s'(i) = s(i) + \epsilon$, with all the other components being the same as in s . Then, $\text{cost}(i, s') - \text{cost}(i, s) = C_i \epsilon - \sum_{\ell: \text{def}(i, \ell) > 0} q_{\ell} p \epsilon = (C_i - p \Pr[\text{def}(i, \ell) > 0]) \epsilon$. Since $s(\cdot)$ is stable, node i has no incentive to increase $s(i)$, which implies that $C \geq p \Pr[\text{def}(i, \ell) > 0]$. Similarly, since node i has no incentive to decrease $s(i)$, we have $C_i \leq p \Pr[\text{def}(i, \ell) > 0]$, and (1) follows from that.

For (2), we have $\text{cost}(i, s) = C_i s(i) - q_{\ell} C_{ij} s(i, j, \ell) + z_1(i, \ell)$ and $\text{cost}(j, s) = C_j s(j) + q_{\ell} C_{ij} s(i, j, \ell) + z_2(j, \ell)$, where $z_1(i, \ell)$ and $z_2(j, \ell)$ are the remaining terms of the respective cost functions. If $s(i, \ell) > D(i, \ell) - s_{in}(i, \ell)$, $\text{def}(j, \ell) > 0$ and $s(i, j, \ell) < \text{cap}(i, j)$ then increasing $s(i, j, \ell)$ (which is possible since $s(i, j, \ell) < \text{cap}(i, j)$) is better for both nodes i and j : (i) this lowers $-q_{\ell} C_{ij} s(i, j, \ell)$ without affecting $\text{def}(i, \ell)$ or other components in $\text{cost}(i, s)$, (ii) since $\text{def}(j, \ell) > 0$ and $C_{ij} \leq p$, this leads to reduction in $\text{cost}(j, s)$, without affecting other components. Therefore part (a) holds. Next, suppose $\text{def}(i, \ell) > 0$ and $C_j \leq p \Pr[\text{def}(j, \ell) > 0]$, decreasing $s(i, j, \ell)$ is better for both i and j : this would reduce $\text{def}(i, \ell)$ thereby lowering $\text{cost}(i, s)$, since $C_{ij} \leq p$. Further, since $C_j \leq p \Pr[\text{def}(j, \ell) > 0]$, node j is better off increasing $s(j)$ and using it to lower the total deficits at node j . Therefore (b) holds. \square

Multiple stable solutions. Stable solutions are not unique for an instance of HSTOCKPILE, in general. This is illustrated in the instance in Figure 2. It can be verified that the strategies $s(\cdot)$ and $s'(\cdot)$ defined in the following manner are both stable: (1) $s(1) = 100$, $s(2) = 0$, $s(3) = 100$, $s(1, 2, 1) = 0$, $s(1, 2, 2) = 100$, $s(3, 2, 1) = 0$, $s(3, 2, 2) = 100$, $s(2, i, \ell) = 0$ for all i, ℓ ; and (2) $s'(1) = 100$, $s'(2) = 100$, $s'(3) = 100$, $s'(1, 2, 1) = 0$, $s'(1, 2, 2) = 50$, $s'(3, 2, 1) = 0$, $s'(3, 2, 2) = 50$, $s'(2, i, \ell) = 0$ for all i, ℓ . We have $\text{cost}(s) = 200$ and $\text{cost}(s') = 300$.

For a graph $G = (V, E)$, a subset $S \subseteq V$ of nodes is a dominating set of G if for each $j \in V$, either $j \in S$ or $N(j) \cap S \neq \emptyset$, where $N(j) = \{i : (i, j) \in E\}$ is the set of neighbors of j . We observe below that in any instance of HSTOCKPILE, the nodes which maintain a stockpile form a dominating set of the graph.

Observation 2. Suppose $C_i < p \Pr_{\ell}[D(i, \ell) > 0]$ for all $i \in V$. Then, for any stable strategy $s(\cdot)$ for an instance $(G, \text{cap}, \mathbf{C}, \mathbf{D}, p)$ of HSTOCKPILE, the set $S = \{i : s(i) > 0\}$ of nodes with positive stockpile forms a dominating set of

G .

Proof. From Lemma 1, it follows that for each node i , $\Pr_{\ell}[\text{def}(i, \ell)] = \frac{C_i}{p} < \Pr_{\ell}[D(i, \ell) > 0]$. This implies that for each node i , there exists at least one scenario ℓ for which $\text{def}(i, \ell) < D(i, \ell)$. This can happen if $s(i, \ell) > 0$ (i.e., node i uses its stockpile) or $s_{in}(i, \ell) > 0$ (i.e., node i purchases from neighbors to reduce its deficits). In other words, each node i either maintains a stockpile, or purchases from a neighbor. Therefore, the set $S = \{i : s(i) > 0\}$ is a dominating set for G . \square

From a policy planner's perspective, the number of nodes which maintain a stockpile is important. Finding stable strategies which minimize this number is hard.

Lemma 3. For an instance $(G, \text{cap}, \mathbf{C}, \mathbf{D}, p)$ of HSTOCKPILE, and parameters k, B determining if there exists a stable solution of cost at most B , in which at most k nodes have stockpiles is NP-complete.

Proof. It is easy to see that the problem is in NP. We show hardness by reduction from the bipartite graph dominating set problem: an instance of this is a bipartite graph $G = (V_1 \cup V_2, E)$ and a parameter k , and the problem is to determine if there exists a subset $S \subset V_1$ with $|S| \leq k$ which dominates V_2 .

We construct an instance of HSTOCKPILE in the following manner: the hospital network is G . We have one scenario $\ell \in \mathcal{D}$ corresponding to each $\ell \in V_2$ with $q_{\ell} = 1/|V_2|$. In scenario ℓ , we have $D(\ell, \ell) = M$, and $D(i, \ell) = 0$ for all nodes $i \in V_1 \cup V_2 - \{\ell\}$. We have $C_i = 0$ for all nodes $i \in V_1$, and $C_i = C$ for all $i \in V_2$. C, p are chosen such that $C < pM/|V_2|$. We have $C_{ij} = C'$ for all $(i, j) \in E$ with $C < C' < p$. We have $\text{cap}(i, j) = M$ for all $(i, j) \in E$.

Suppose there exists a stable solution $s(\cdot)$ in this game instance of cost 0, with stockpile at at most k nodes. In order to have total cost 0, the nodes with stockpiles must all be in V_1 , and all nodes must have all deficits 0. Let $S \subset V_1$ denote the set of nodes with positive stockpile. Then, $|S| \leq k$. Since all nodes in V_2 have 0 deficits, for each $\ell \in V_2$, there exists $i \in V_1$ such that $s(i, \ell) > 0$. Therefore, S is a dominating set for V_2 of size at most k .

Next, suppose there exists a dominating set $S \subset V_1$ of size at most k . We set $s(i) = |N(i)|M$. For each $\ell \in V_2$, there exists $i \in S$ such that $(i, \ell) \in E$. We set $s(i, \ell) = M$. Then, $s(\cdot)$ is a stable strategy—no node $i \in V_1$ has incentive benefits by reducing or increasing its stockpile. Since $C_{i\ell} < p$, node $\ell \in V_2$ does not benefit by reducing the transfer, and has no incentive to increase it, since its deficit is already 0. Therefore $s(\cdot)$ is a stable strategy, and has cost 0. \square

Lemma 4. For any instance $(G, \text{cap}, \mathbf{C}, \mathbf{D}, p)$ of HSTOCKPILE in which $C_i = C$ for all $i \in V$, the price of anarchy is bounded by $(\max_{i, \ell} D(i, \ell))(\sum_{(i, j) \in E} \text{cap}(i, j))$. In the absence of capacity constraints, there exist instances of HSTOCKPILE, for which the price of anarchy is unbounded.

Proof. (Sketch) Let $s^*(\cdot)$ denote the strategy with the minimum cost. Let \mathcal{I}_j denote restricted to node j —this consists of a graph with the single node j , with stockpiling cost

$C_j = C$, and demands $D(j, \ell)$ with probability q_ℓ . Let OPT_i denote the minimum cost of any strategy for \mathcal{I}_j . Then, $\text{cost}(s^*) \geq \max_j OPT_j$. Next, consider any stable strategy $s(\cdot)$. For any node j , $s_{out}(j, \ell) \leq \sum_{(ij) \in E} \text{cap}(i, j)$, and $\text{cost}(j, s) \leq C \max_{j, \ell} D(j, \ell)$.

Next, suppose there are no capacity constraints. Consider a star graph G on n nodes $\{1, 2, \dots, n\}$, with node 1 being the center of the star. We have \mathcal{D} with $n - 1$ scenarios, each corresponding to node $j = 2, \dots, n - 1$. Let $D(\ell, \ell) = M$ for $\ell = 2, \dots, n$ and $D(j, \ell) = 0$ for all $j \neq \ell$. Also, $q_\ell = \frac{1}{n-1}$. Let $s(\cdot)$ be a strategy vector such that $s(j) = 0$ for all $j > 1$, $s(1) > (n - 1)M$, and $s(1, \ell, \ell) = s(1)$. We have $C_i = C$ and $p \geq C(n - 1)$. Then, $\text{cost}(s) = s(1)C$. Also, $s(\cdot)$ is stable, since node 1 has negative cost, and therefore, has no incentive to reduce $s(1, j, \ell)$ for any j, ℓ . On the other hand, each node $j > 1$ has no incentive to increase $s(1, j, \ell)$. Next, consider a strategy vector $s'(\cdot)$ that has $s'(j) = M$ for all $j = 2, \dots, n - 1$, $s'(1) = 0$, and $s'(i, j, \ell) = 0$ for all i, j . Then, $s'(\cdot)$ is also a stable solution because $C \leq p/(n - 1)$, so that no node $j > 1$ has incentive to reduce $s'(j)$. Since this leads to zero deficit, no node has incentive to increase its stockpile either. We have $\text{cost}(s') = (n - 1)M$, so that the price of anarchy is at least $\text{cost}(s)/\text{cost}(s') > \frac{s(1)}{M(n-1)}$, which is unbounded. \square

Given a strategy $s(\cdot)$, we say that it is a local optimum with respect to any changes in $s(i)$, $s(j)$ and $s(i, j, \ell)$ for any i, j, ℓ if $\text{cost}(s) \leq \text{cost}(s')$ for any strategy $s'(\cdot)$ such that: (a) $s'(k) = s(k)$, for all $k' \neq i, j$, (b) for all i', j' such that $\{|i', j'\} \cap \{i, j\}\} \leq 1$, $s'(i', j', \ell) = s(i, j, \ell)$.

Theorem 5. Let $s(\cdot)$ be a strategy profile such that $\text{cost}(s)$ is a local optimum with respect to any changes in $s(i)$, $s(j)$ and $s(i, j, \ell)$ for any choice of i, j, ℓ . Then, $s(\cdot)$ is a stable strategy. Conversely, if $s(\cdot)$ is stable, it is a local optimum.

Proof. We show that $s(\cdot)$ satisfies all the conditions of a stable solution. Consider a node i . Let $s'(\cdot)$ be a strategy that agrees with $s(\cdot)$ in all components, except possibly $s(i)$. Then, $\text{cost}(i', s') = \text{cost}(i', s)$ for all $i' \neq i$, since $s'(i') = s(i')$ and $s'(i', j, \ell) = s(i', j, \ell)$ for all j, ℓ , by definition of $s'(\cdot)$. Since $s(\cdot)$ is a local optimum, it follows that $\text{cost}(s) = \sum_j \text{cost}(j, s) \leq \text{cost}(s') = \sum_j \text{cost}(j, s')$, which implies that $\text{cost}(i, s) \leq \text{cost}(i, s')$. Therefore, no node i has incentive to unilaterally deviate.

Next, consider a pair of nodes i, j , and a strategy $s'(\cdot)$ that possibly differs from $s(\cdot)$ in $s(i)$, $s(j)$, $s(i, j, \ell)$, but is identical to $s(\cdot)$ in all other components. Then, for $i' \neq i, j$, $\text{cost}(i', s') = \text{cost}(i', s)$ because $s'(i') = s(i')$ and $s'(i', j', \ell) = s(i', j', \ell)$, unless $\{i', j'\} = \{i, j\}$. Since $s(\cdot)$ is a local optimum, we have $\text{cost}(s) \leq \text{cost}(s')$, which in turn implies that $\text{cost}(i, s) + \text{cost}(j, s) \leq \text{cost}(i, s') + \text{cost}(j, s')$. It follows that either (a) $\text{cost}(i, s) \leq \text{cost}(i, s')$ and $\text{cost}(j, s) \leq \text{cost}(j, s')$ or (b) $\text{cost}(i, s) < \text{cost}(i, s')$ or $\text{cost}(j, s) < \text{cost}(j, s')$. This implies that no pair of nodes i, j have incentive to deviate. Therefore $s(\cdot)$ is a stable strategy.

Next, suppose $s(\cdot)$ is stable. Consider any $s'(\cdot)$ that differs from $s(\cdot)$ in $s(i)$, $s(j)$, $s(i, j, \ell)$, (j, i, ℓ) , but agrees on all other components. This implies $\text{cost}(i', s) = \text{cost}(i', s')$ for

all $i' \neq i, j$. Since $s(\cdot)$ is stable, $\text{cost}(i, s) + \text{cost}(j, s) \leq \text{cost}(i, s') + \text{cost}(j, s')$. Together, this implies $\text{cost}(s) \leq \text{cost}(s')$, so that $s(\cdot)$ is a local optimum. \square

This implies that any local optimum is a stable strategy. In particular, the optimum is also stable.

Corollary 6. Let $s^*(\cdot)$ be a solution with the minimum possible cost. Then, $s^*(\cdot)$ is stable.

5 The social optimum and reducing the deficit

We now show that optimal strategies $s(\cdot)$ can be computed by linear programming.

$$\begin{aligned} & \min \sum_i C_i s(i) + \sum_{i, \ell} q_\ell p f(i, \ell) \text{ s. t.} \\ & s(i) = s(i, \ell) + \sum_{j \neq i} s(i, j, \ell), \quad \text{for all } i, \ell \\ & f(i, \ell) \geq D(i, \ell) - s(i, \ell) - \sum_{j \neq i} s(j, i, \ell) \\ & \quad + \sum_{j \neq i} s(i, j, \ell), \quad \text{for all } i, \ell \\ & f(i, \ell) \geq 0, \quad \text{for all } i, \ell \\ & s(i) \geq 0, \quad \text{for all } i \\ & s(i, \ell) \geq 0, \quad \text{for all } i, \ell \\ & s(i, j, \ell) \leq \text{cap}(i, j), \quad \text{for all } i \neq j, \ell \\ & s(i, j, \ell) \geq 0, \quad \text{for all } i \neq j, \ell \end{aligned}$$

Lemma 7. The solution $s^*(\cdot)$ computed by the above linear program has the minimum cost.

Proof. From the constraints of the above program, it follows that $f(i, \ell) \geq \max\{D(i, \ell) - s(i, \ell) - \sum_{j \neq i} s(j, i, \ell) + \sum_{j \neq i} s(i, j, \ell), 0\}$. Since the objective is to minimize $\sum_i C_i s(i) + \sum_{i, \ell} q_\ell p f(i, \ell)$, it follows that for each i, ℓ , $f(i, \ell)$ will actually satisfy the above inequality by an equality. Next, the lower and upper bounds ensure that all $s(i, j, \ell)$ are feasible.

Recall that $\text{cost}(i, s) = C_i s(i) + \sum_\ell \sum_{j \neq i} C_{ji} s(j, i, \ell) - \sum_\ell \sum_{j \neq i} C_{ij} s(i, j, \ell) + \sum_\ell q_\ell p \text{def}(i, \ell)$. Therefore, the components $C_{ij} s(i, j, \ell)$ contribute a positive term to $\text{cost}(i, s)$, and a negative term to $\text{cost}(j, s)$. These cancel out in $\sum_i \text{cost}(i, s)$, and is precisely equal to the objective function of the above program. Therefore, $s^*(\cdot)$ is a solution that minimizes $\sum_i \text{cost}(i, s)$. Finally, from Lemma 6, it follows that s^* is a stable solution, and the lemma follows. \square

Reducing deficits for a given strategy

We now consider the problem MINDEF of minimizing the total deficit, given a specific strategy $s(\cdot)$, from a centralized agency's perspective. Formally, this problem is defined in the following manner: Given an instance (G, cap, C, D, p) of HSTOCKPILE, a specific strategy $s(\cdot)$, budget B and a bound




Figure 3: A hospital network induced by some of the largest hospitals in North Carolina.

k , the objective of the $\text{MINDEF}(s, B, k)$ problem is to select an additional stockpile vector \hat{s} , such that

1. $\sum_i \hat{s}_i \leq B$ and $|\{i : \hat{s}_i > 0\}| \leq k$, which captures the constraint that a centralized agency can only send the supplies to a bounded number of hospital nodes.
2. $\sum_i \text{def}(i, s + \hat{s})$ is minimized.

The MINDEF problem turns out to be very hard in general, as discussed below.

Lemma 8. *The $\text{MINDEF}(s, B, k)$ problem is NP-hard.*

6 Experimental Results

We study the HSTOCKPILE formulation on a hospital network formed by the ten largest hospitals in the state of North Carolina, as shown in Figure 3. All of these hospitals are Level I or Level II trauma centers that would reasonably be expected to provide patient care during a public health emergency. We assume a completely connected undirected network between these ten hospitals, with the cost of sharing supplies between hospitals dictated solely by the distance between them. This, in effect, creates a network with several clusters of hospitals that can share easily—those within the “Triangle” made up of the cities of Raleigh, Durham and Chapel Hill, as well as those in the cities of Winston-Salem and Greensboro—as well as more distant hospitals who will have a more difficult time obtaining supplies from other hospitals. We choose the cost C_{ij} to be proportional to the distance $d(i, j)$ between the hospitals, and choose $C_i = \min_{i,j} d(i, j)$. We vary the penalty p .

Epidemics are inherently stochastic and unpredictable, making the allocation of supplies a difficult problem. Stockpiling strategies must be robust to both unexpectedly severe outbreaks, such as the 2009 H1N1 influenza epidemic, as well as less severe outbreaks that may not necessitate extensive stockpiles. To simulate this varying demand, we simulate 1000 stochastic epidemics of a flu-like illness, using a variation of a Susceptible-Exposed-Infected-Recovered (SEIR) epidemic model, where infected patients are subdivided into ten compartments (I_j), indicating that the demand for their care is the responsibility of Hospital j , allocated proportionally based on a weight (κ_j). In one experiment, this weight was proportional to the bed-size of the hospital, and in another, the weight was assigned randomly.

$$\begin{aligned} \frac{dS}{dt} &= -\frac{\beta S \sum_{j=1}^{10} I_j}{N} \\ \frac{dE}{dt} &= \frac{\beta S \sum_{j=1}^{10} I_j}{N} - \alpha E \\ \frac{dI_j}{dt} &= \alpha \kappa_j E - \gamma I_j \\ \frac{dR}{dt} &= \gamma \sum_{j=1}^{10} I_j \\ \frac{dN}{dt} &= S + E + \sum_{j=1}^{10} I_j + R \end{aligned}$$

We choose the parameters so that they correspond to pandemic-grade flu, with the reproductive number R_0 (which corresponds to the expected number of secondary infections caused by any individual) between 1.7 and 2.0 (Halloran et al. 2008).

Figure 4 shows the stockpile levels for the social optimum solution computed using the algorithm from Section 5 for the ratio C/p varying from 0.5 to 1. The plots show the stockpile levels and deficits for each of the ten hospitals. As the plots show, at low penalties, the stockpile levels are much lower than the expected deficits at individual hospitals. This is flipped when the penalty p becomes much higher. We expect a much higher variability for real datasets.

7 Conclusions and acknowledgements

We formalize the hospital stockpiling game problem (HSTOCKPILE) with network capacity constraints as a non-cooperative game, using the notion of pairwise Nash stability. We show that stable solutions might still have fairly high levels of deficits, and this motivates a deeper study of incentives to reduce the deficits, including centralized stockpiling, and changing penalties. The network structure has a significant impact on these problems, and makes these problems computationally challenging. Our results suggest that other notions of stability, which consider more general incentives for pairs of nodes to deviate, might be needed for studying HSTOCKPILE . Extending our approach to incorporate more realistic constraints are important open problems. This work has been partially supported by the following grants: DTRA Grant HDTRA1-11-1-0016, DTRA CNIMS Contract HDTRA1-11-D-0016-0010, NSF ICES CCF-1216000, NSF NETSE Grant CNS-1011769, NIH MIDAS Grant 5U01GM070694, NSF DIBBS Grant ACI-1443054.

References

- Aspnes, J.; Chang, K.; and Yampolskiy, A. 2006. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *J. Comput. Syst. Sci.*
- Bauch, C., and Earn, D. 2004. Vaccination and the theory of game. *PNAS*.
- Blume, L.; Easley, D.; Kleinberg, J.; Kleinberg, R.; and Tardos, E. 2012. Network formation in the presence of contagious risk. In *Proc. IEEE FOCS*.
- DeLaurentis, P.-C.; Adida, E.; and Lawley, M. 2011. Hospital stockpiling for disaster planning. *IIE Transactions* 348–362.

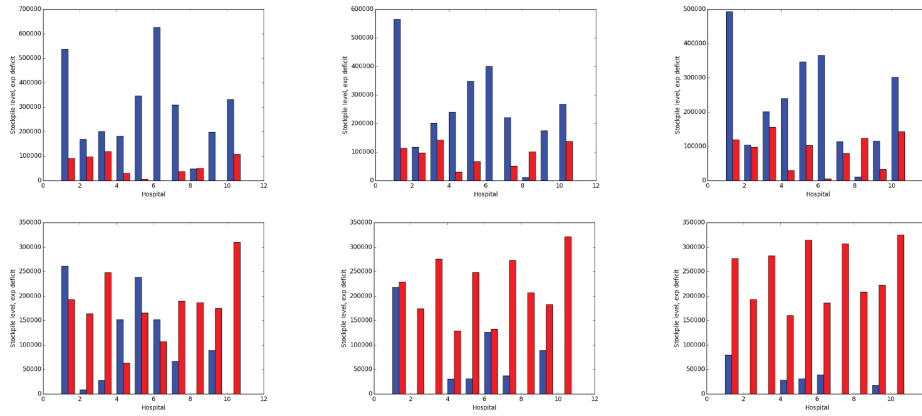


Figure 4: Stockpile level (blue bars) and expected deficit in demand (red bars) for each hospital, for the C/p ratio equal to 0.5, 0.6, 0.7, 0.8, 0.9 and 1.0, respectively.

DeLaurentis, P.-C. 2009. Hospital stockpiling for influenza pandemics. Ph. D thesis, Purdue University.

Donker, T.; Wallinga, J.; and Grundmann, H. 2010. Patient referral patterns and the spread of hospital-acquired infections through national health care networks. *PLoS Computational Biology*.

Halloran, M. E.; Ferguson, N. M.; Eubank, S.; Longini, I. M.; Cummings, D. A.; Lewis, B.; Xu, S.; Fraser, C.; Vullikanti, A.; Germann, T. C.; et al. 2008. Modeling targeted layered containment of an influenza pandemic in the united states. *Proceedings of the National Academy of Sciences* 105(12):4639–4644.

Jackson, M. 2008. *Social and Economic Networks*. Princeton University Press.

Koutsoupias, E., and Papadimitriou, C. H. 1999. Worst-case equilibria. In *Proceedings of STACS*.

Kumar, V. S. A.; Rajaraman, R.; Sun, Z.; and Sundaram, R. 2010. Existence theorems and approximation algorithms for generalized network security games. In *Proceedings of IEEE ICDCS*.

Lee, B. Y.; McGlone, S. M.; Yilmaz, K. F. W. S. L.; Avery, T. R.; Song, Y.; Christie, R.; Eubank, S.; Brown, S. T.; Epstein, J. M.; Parker, J. I.; Burke, D. S.; Platt, R.; and Huang, S. S. 2011. Modeling the spread of methicillin-resistant staphylococcus aureus (mrsa) outbreaks throughout the hospitals in orange county, california. *Infect Control Hosp Epidemiol*. 562–572.

Leyton-Brown, K., and Shoham, Y. 2008. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers, 1st edition.

Marathe, M., and Vullikanti, A. 2013. Computational epidemiology. *Communications of the ACM*.

Meca, A.; Timmer, J.; Garcia-Jurado, I.; and Borm, P. 2004. Inventory games. *European Journal of Operational Research* 127–139.

Meyers, L. A., and Dimitrov, N. 2010. Mathematical ap-

proaches to infectious disease prediction and control. *INFORMS, Tutorials in Operations Research*.

Omic, J.; Orda, A.; and Miegheem, P. V. 2009. Protecting against network infections a game theoretic perspective. In *INFOCOM*.

Prakash, B. A.; Adamic, L. A.; Iwashyna, T. J.; Tong, H.; and Faloutsos, C. 2013. Fractional immunization in networks. In *Proc. of SDM*, 659–667.

Reluga, T., and Galvani, A. 2011. A general approach to population games with application to vaccination. *Mathematical Biosciences*.

Saha, S.; Adiga, A.; and Vullikanti, A. K. S. 2014. Equilibria in epidemic containment games. In *The 28th AAAI Conference on Artificial Intelligence (AAAI)*.

Simmering, J. E.; Polgreen, L. A.; Campbell, D. R.; Cavanaugh, J. E.; and Polgreen, P. M. 2015. Hospital transfer network structure as a risk factor for clostridium difficile infection. *Infect Control Hosp Epidemiol*.

Stier, D. D., and Goodman, R. A. 2007. Mutual aid agreements: Essential legal tools for public health preparedness and response. *American Journal of Public Health* S62–S68.

Trajanovski, S.; Hayel, Y.; Altman, E.; Wang, H.; and Van Mieghem, P. 2015. Decentralized protection strategies against sis epidemics in networks. *Control of Network Systems, IEEE Transactions on PP(99)*:1–1.