

Abstract

In response to the increasing threat of terrorist attacks and natural disasters, governmental and private organizations world-wide have invested significant resources in disaster planning activities. This paper addresses joint inventory stockpiling of medical supplies for groups of hospitals prior to a disaster. Specifically, we consider the problem of determining the stockpile quantity of a medical item at several hospitals. We assume demand is uncertain and driven by the characteristics of a variety of scenarios of a disastrous event. Further, we assume hospitals have mutual aid agreements for inventory sharing in the event of a disaster. Each hospital's desire to minimize its stockpiling cost together with the potential to borrow from other stockpiles creates individual incentives well represented in a game theoretic framework. We model this problem as a non-cooperative 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. The comparison provides some managerial insights and public health policy implications valuable for disaster planning.

1 Introduction

The U.S. healthcare system faces challenges from disastrous events. A timely example is the recent outbreak of H1N1 influenza pandemic. In the U.S., the reported percentage of visits for influenza-like illness (ILI) of Week 40 (ending October 10, 2009) is 6.1% of all visits, almost 3 times higher than the national baseline of 2.3% (The Centers for Disease Control and Prevention (CDC), 2009). This translates into an approximate increase of 42.7 million patients presenting ILI at healthcare facilities across the country (this number is estimated based on the 2006 total number of visits to physician offices, hospital outpatient departments, and hospital emergency departments, published by The Centers for Disease Control and Prevention (CDC) 2008). Influenza-associated hospitalizations have also increased tremendously over previous seasons. This shows how adverse events such as a flu pandemic can create a significant surge on the healthcare system.

Public health organizations and healthcare industries have intensified their disaster preparedness and response efforts in recent years as a result of the increasing threat of terrorist attacks and natural hazards. These planning efforts include the establishment of disaster response procedures and protocols for all hazards, and cooperative relationships between government agencies and healthcare providers. Some governmental entities, such as the World Health Organization (WHO), the Department of Health and Human Services (HHS), and CDC provide general guidelines and recommendations for setting and achieving preparedness goals such as listed on <http://Flu.gov/>. However, these guidelines typically lack the detailed instructions and/or methods that 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. Unfortunately, government guidelines do not address the financial implication of stockpiling supplies to healthcare providers. In addition, as Toner and Waldhorn (2006) 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. It is also important to point out that stockpiling large quantities of supplies is a financial challenge for hospitals. The Center for Biosecurity estimated that a 164-bed hospital would need to spend at least \$640,000 to stockpile minimal personal protective equipment (PPE, such as masks, gloves, gowns, etc.) and basic

supplies for a 1918-like pandemic (Toner and Waldhorn, 2006). In addition, hospitals will need to absorb extra costs for stock rotation and staff education/training as well as proper arrangement of sufficient ancillary services (such as waste and laundry) required by additional use of PPEs, all of which further impacts their bottom line. Therefore, 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 (Havlak et al., 2002).

Our work focuses on stockpile decision making within a community of hospitals. We adopt a game theoretical framework to account for the interactions of hospitals located in the same area. There are several types of such interactions. First, these hospitals may be located close enough to each other to serve overlapping populations, and hence patients may have a choice of which hospital to go to in an emergency. As a result, demand can be re-distributed from one hospital to another if the former experiences a shortage. Moreover, hospitals in the same area commonly engage in mutual aid agreements, such as memoranda of understanding (MOU), to share supplies with those encountering shortages during emergency situations. Because of these interactions, hospital stockpile decisions are inter-dependent: the stockpiling decision of one hospital in the group affects others in the group, which places the problem in a game theoretical setting. Note, however, that even though decisions made by hospitals in the group have mutual impact, each hospital is a distinct decision-maker motivated by its own objectives in preparing emergency plans. In our experience, there is often little or no cooperation among hospitals in the disaster preparedness planning stage, particularly in areas where competition between healthcare organizations is intense. The MOUs simply state some cooperative agreement on supplies sharing *after* a disaster happens. Therefore, we think that a non-cooperative game is the appropriate framework for this problem.

We assume that hospitals want to meet as much of the surge demand in a disaster as possible, at the lowest possible cost, considering that 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 in each scenario. We further assume that in the case of a system-wide shortage 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, potentially contributing to further supply shortages at other hospitals.

We note that our approach uses a stylized game theoretical model. Such models are never strict representations of reality; for example, our models do not include service level or space constraints. However, such models make analytical solution possible and thus facilitate the discovery of fundamental behaviors and insights that would otherwise remain hidden. This might not be possible with more complex representations. The model presented here encompasses the major features of the problem and serves the purpose of yielding interesting insights while remaining tractable.

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? The stockpile of supplies considered in this problem is the amount beyond what the hospital needs under normal conditions, and the surge patient demand during a disaster is expected to be much larger than what the hospital experiences during regular operations. We assume that at the time that the stockpile decisions considered in this paper are made, the current stockpile is zero (or negligible). Moreover, we neglect safety stocks or functional inventory of supplies stored by the hospital due to normal demand variations for the reason that some of these supplies may still be needed for treating non-disaster victims. We also assume non-zero lead time before newly ordered supplies can be provided by suppliers since production ramp-up is unlikely to be fast. As a result, when the disaster occurs, there is a phase during which the only supplies available are those that have been previously stockpiled. This phase ends when suppliers make new deliveries. The demand considered in our model corresponds to the demand surge due to the disaster during this phase, i.e. the demand that can only be satisfied using the available stockpiles (Julian Fantino, 2003; Hung, 2003; Chan-Yeung, 2004).

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 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.

1.1 Relevant literature

The hospital stockpiling problem is related to the supply chain management and inventory control research. In particular, the models of inventory with transshipment are worth noting; some representative work includes Krishnan and Rao (1965), Tagaras (1989), and Rudi et al. (2001). Zhao et al. (2005) consider inventory sharing in a decentralized dealer network focusing on *emergency lateral transshipment* which occurs only when one of the participating entities has shortage and thus needs transshipment from other entities. They examine the base-stock, rationing and sharing policy as decision variables in an inventory sharing and rationing game model. In contrast, our hospital stockpiling game model considers each hospital's stockpile level as its decision variable with consideration of several possible demand scenarios based on disaster risk assessment, with supply sharing policy fixed by the mutual aid agreement in place, or by the automatic nature of demand redistribution. Another contrast in our model is the penalty imposed on each hospital should a community-wide supply shortage occur for the reasons described earlier, while Zhao et al. (2005) consider lost sales, backorder, and delay costs as individual dealers fail to fill customer demand. More fundamentally, in our model hospitals are not only trying to minimize financial cost, but also taking into account their mission of serving the population and providing life-saving supplies, while Zhao et al. (2005) as well as most of the literature on inventory transshipment takes a purely profit-driven approach.

Mathematical modeling has been applied to medical stockpiling problems. One class is the application of inventory modeling which is commonly seen in the field of supply chain management. Jacobson et al. (2006) use a stochastic inventory model to analyze whether or not the CDC-proposed pediatric vaccine stockpile levels are sufficient, with the consideration of vaccine production interruptions. Another class of mathematical modeling is the application of cost-benefit analysis that typically weighs available options from economic viewpoints. Lee et al. (2006) examine three neuraminidase inhibitor (e.g. Oseltamivir) stockpiling strategies in responding to a flu pandemic in Singapore: no action, treatment only, and prophylaxis. The results show that the treatment only option is most economical, but implementing a prophylaxis strategy would save the most lives.

Moreover, Medema et al. (2004) present a case study involving a computer-based simulation model which combines a vaccine production model and a cost-effectiveness model of flu intervention strategies. The analysis shows that vaccination strategies would be the most cost-effective for the elderly populations in three European countries.

Recently, there have been several game theoretic models applied to the healthcare arena. Chick et al. (2008) study contractual issues of the influenza vaccine supply chain as a theoretic game between a manufacturer and the government. Their analysis shows that a rational manufacturer will always under-produce vaccines due to the risk of uncertain vaccine production processes. Sun et al. (2009) study the problem of antiviral and vaccine stockpiling in different countries and possible international resource sharing should an epidemic start in a country possessing little or no such supplies. They combine an epidemic transmission model and a game theoretic setup to examine drug allocations among participating countries. Wang et al. (2009) study a game between selfish countries that each allocate resources at the onset of an epidemic to minimize the total number of infected citizens. They show that selfish countries will allocate resources to lower the disease reproductive ratio below a threshold in order to avoid a major outbreak.

1.2 Organization

This research focuses on stockpiling medical supplies in a group of hospitals as a disaster preparedness measure. It includes several contributions to the current literature and practice. To our knowledge, this is the first paper to model and analyze disaster stockpiling for a community of hospitals sharing supplies or serving a common population. Our model captures uncertainty by considering a variety of disastrous scenarios. It captures the fact that hospitals share excess supplies in a medical emergency based on their mutual aid relationships or due to overlapping populations. In addition, the research integrates these and other relevant factors into a game theoretic setting, develops closed-form solutions for Nash equilibria and the centralized solution. We compare centralized and decentralized stockpiling strategies of a group of hospitals to provide public health management insights.

The organization of the paper is as follows. Section 2 describes the hospital stockpiling game model (i.e. the decentralized setting), its properties, and the Nash solutions. Section 3 discusses the stockpiling problem with a centralized coordination. Section 4 compares the Nash and decentralized

solutions and draws insights. Finally, Section 5 concludes the work.

2 Hospital Stockpiling Game Model

2.1 Model setup

We consider a group of hospitals $\{1, \dots, n\}$ serving a given region. Each hospital i needs to decide the stockpile level s_i of its supply for a medical item in anticipation of a disaster. In this model, the severity and type of disaster is uncertain. As a result, the overall patient demand, D , for the item in the region of concern is also uncertain. We consider several disaster scenarios, $l = 1, \dots, L$ with probability q^1, \dots, 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 given that scenario l occurs), with a cumulative distribution function, $F^l(\cdot)$ differentiable on \mathbb{R} , and a probability density function, $f^l(\cdot)$ defined on \mathbb{R} and strictly positive on $[a_l, b_l] \subset \mathbb{R}^+$ (b_l may be infinite). Note that D^l is the random variable that represents the total aggregate demand to be served by the entire group of hospitals in scenario l . Since in our model, only the system-wide shortage matters, we make no assumption on the stochasticity of demand at individual hospitals. In particular, the different coefficients of variation (ratio of standard deviation and mean) of the demand at different hospitals do not play a role.

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 supplies 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.

The modeling assumption of a system-wide distributed penalty caused by a system-wide supply shortage, rather than individual penalties due to localized shortages, is a critical assumption in this paper. In particular, it implies that while the stockpiling decision made by a hospital is directly connected to its stockpiling cost, it does not directly affect its utility: instead, it affects it *indirectly* via the total quantity stockpile by all hospitals in the group. It can be justified by one of three

externality factors that are specific to the context we consider of disaster planning for hospital in a common area: (1) increased risk of disease propagation in the community (for the case of a disaster related to an infectious disease), (2) demand redistribution, (3) sharing of resources. First, when one hospital experiences a shortage and is unable to provide appropriate treatment or prophylaxis to protect its patients against an infectious disease, the disease will propagate faster in the entire community served by the group of hospitals, and more people will get infected not only in the fraction of the population served by the hospital experiencing a shortage, but also in entire region served by the group of hospitals, which will affect other hospitals. Second, in large metropolitan areas served by multiple healthcare facilities, patients may easily switch from an under-supplied facility to one with more supplies. Therefore, when a given hospital experiences a shortage, the population that is now unable to receive the necessary supplies from this hospital is likely to turn to other hospitals in the area, and thus these hospitals that would otherwise have enough supplies may incur a shortage as well due to the increased demand. A third justification for making this assumption is the sharing of resources that occurs in practice based on MOUs that often exist between hospitals of a same area. These MOUs state that in case of an emergency, hospitals agree to help each other if in need. In particular, when one hospital is short of supplies, it may ask other hospitals to share some of their supplies. As a result, other hospitals in the group are negatively affected by the shortage at a given hospital and will observe a higher demand for their own limited resources and/or will have to share their own supplies. This implies what matters at the time of a disaster in the region is not only how much supply a hospital has available, but how prepared every hospital in the region is, i.e. each hospital is affected not by its own stockpile but by the cumulative stockpile (and shortage). A shortage in the overall system will hurt not one but all of the hospitals involved.

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. While we recognize that selecting a meaningful value of p is critical, the task is not within the scope of this study. We will provide some sensitivity analysis on the value of p to show its impact on the solutions in this report.

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 represents the relative stake of hospital i in the overall supplies shortage. Its estimation must reflect the relative effect on hospital i of a system-wide shortage. There is a variety of ways to evaluate this parameter. One way could be to interpret the fraction w_i as hospital i 's fraction of service capacity to the overall system capacity (typically estimated by its bed size); another way could be to interpret it as hospital i 's share of the responsibility in the region's preparedness.

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

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

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 $\mathbf{s} = (s_1, \dots, s_n)$ through the dependency of U on S .

The following is an important assumption of this model.

Assumption 1. $c_i < p w_i$, $i = 1, \dots, n$.

The interpretation of this assumption is as follows. 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.

2.2 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

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

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

$$\min_{s_i \geq 0} J_i(s_i, \mathbf{s}_{-i}) = \min_{s_i \geq 0} c_i s_i + p w_i \sum_{l=1}^L q^l \int_S^\infty (x - S) f^l(x) dx.$$

We now analyze the characteristics of hospital i 's cost function J_i .

Proposition 1. $J_i(s_i, \mathbf{s}_{-i})$ is a continuous, twice differentiable, convex function of s_i .

Proof. Continuity and differentiability follow directly from the definition of $J_i(s_i, \mathbf{s}_{-i})$. To show convexity, we calculate the first and second partial derivatives.

$$\begin{aligned} \frac{\partial J_i(s_i, \mathbf{s}_{-i})}{\partial s_i} &= c_i - p w_i \sum_{l=1}^L q^l \int_S^\infty f^l(x) dx \\ &= c_i - p w_i \sum_{l=1}^L q^l [1 - F^l(S)] \\ &= c_i - p w_i + p w_i \sum_{l=1}^L q^l F^l(S) \end{aligned} \tag{2}$$

where the last equality follows from $\sum_{l=1}^L q^l = 1$.

$$\frac{\partial^2 J_i(s_i, \mathbf{s}_{-i})}{\partial s_i^2} = p w_i \sum_{l=1}^L q^l f^l(S) \geq 0.$$

□

The following proposition gives the best response function in a closed-form. Let $\bar{G}(s) = \sum_{l=1}^L q^l F^l(s)$ 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 $G(\cdot)$ the restriction of $\bar{G}(\cdot)$ to $[a, b]$, thus $G(\cdot)$ is strictly increasing and invertible from $[a, b]$ into $[0, 1]$, and $G^{-1}(\cdot)$ is strictly increasing from $[0, 1]$ into $[a, b]$.

Proposition 2. The best response function, i.e., the solution to the best response problem, can be

written as follows.

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

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 .

Proof. Using the first order necessary conditions, (2) and the fact that $S = s_i + S_{-i}$, the best response s_i is such that

$$c_i - pw_i + pw_i \bar{G}(s_i + S_{-i}) = 0,$$

if this solution is non-negative. The result follows. \square

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




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

observe that there is 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.

2.3 Nash equilibrium solution

Inventory levels $(\bar{s}_1, \dots, \bar{s}_n)$ form a Nash equilibrium if no firm can decrease its expected cost by unilaterally altering its stockpile level:

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

Graphically, a Nash equilibrium is any point that lies at the intersection of the best response functions of all the players in the game.

The following assumption is made without loss of generality.

Assumption 2. *The hospitals are ordered so that*

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

where $r_i = \frac{c_i}{w_i}$ and m is defined as the total number of hospitals that have the minimum ratio r_i in the group.

We note that by Assumption 1, $0 < 1 - r_i/p (< 1)$. We now find the Nash equilibrium solution in closed-form.

Theorem 1. *If $m > 1$, there are infinitely many Nash equilibria given by*

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

for any $\alpha_1, \dots, \alpha_m \in [0, 1]$ such that $\sum_{i=1}^m \alpha_i = 1$.

Proof. It is easy to see that the solution provided satisfies $\bar{s}_i = s_i^*(\bar{\mathbf{s}}_{-i})$, $i = 1, \dots, n$. To show that this is the only solution, let $\bar{s}_1, \dots, \bar{s}_n$ an equilibrium solution, i.e. $\bar{s}_i = s_i^*(\bar{\mathbf{s}}_{-i})$, $i = 1, \dots, n$, $\bar{S} = \sum_{i=1}^n \bar{s}_i$ and $\bar{S}_{-i} = \sum_{j \neq i} \bar{s}_j$. If $\bar{s}_i = 0 \forall i$, then $\bar{S}_{-i} = 0 \forall i$, so $G^{-1}(1 - \frac{r_i}{p}) > S_{-i}$ and thus from Proposition 2, $\bar{s}_i = G^{-1}(1 - \frac{r_i}{p}) > 0$, which is a contradiction. Therefore at least one of the \bar{s}_i is positive, say \bar{s}_{i_0} . It follows from the best response function at hospital i_0 that

$\bar{S} = \bar{s}_{i_0} + \bar{S}_{-i_0} = G^{-1}(1 - \frac{r_{i_0}}{p})$ so all stockpiles at equilibrium must add up to the value $G^{-1}(1 - \frac{r_{i_0}}{p})$. For any i such that $\bar{s}_i > 0$, it follows from the best response function at hospital i that $\bar{S} = G^{-1}(1 - \frac{r_i}{p}) = G^{-1}(1 - \frac{r_{i_0}}{p})$ so $r_i = r_{i_0}$. For any i such that $\bar{s}_i = 0$, $\bar{S}_{-i} = \bar{S} - \bar{s}_i = \bar{S} = G^{-1}(1 - \frac{r_{i_0}}{p})$ and thus, from the best response function at hospital i it follows that $G^{-1}(1 - \frac{r_i}{p}) < G^{-1}(1 - \frac{r_{i_0}}{p})$. Because G^{-1} is strictly increasing, it follows that $r_i > r_{i_0}$. \square

Corollary 1. *If $m = 1$, there is a unique Nash equilibrium given by*

$$\bar{s}_i = \begin{cases} G^{-1}(1 - \frac{r_i}{p}) & i = 1 \\ 0 & i = 2, \dots, n. \end{cases}$$

The proof is similar to the proof of Theorem 1 for $m = 1$, and is thus omitted.

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.

The closed-form expression for the total stockpile at equilibrium follows directly from the above results.

Corollary 2. *The system-wide total stockpile level at any equilibrium, S_d , is given by*

$$S_d = \sum_{i=1}^n \bar{s}_i = G^{-1}\left(1 - \frac{r_1}{p}\right), \quad \forall m, \alpha_1, \dots, \alpha_m.$$

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 , $i = 1, \dots, n$, which is Q_1 .

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

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

We note in particular that Π_d^* depends on $\alpha_1, \dots, \alpha_m$ 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 $c^{max} = \max\{c_1, \dots, c_m\}$. 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.

Figure 2 shows the Nash equilibrium solutions in the case of two hospitals. When the hospitals have different ratios of $r_1 = \frac{c_1}{w_1}$ and $r_2 = \frac{c_2}{w_2}$, there is only one Nash equilibrium as shown on the left-hand side graph. Assuming $\frac{c_1}{w_1} < \frac{c_2}{w_2}$, at the equilibrium hospital 1 stockpiles the amount $Q_1 = G^{-1}\left(1 - \frac{r_1}{p}\right)$ and hospital 2 stockpiles nothing. On the other hand, when $\frac{c_1}{w_1} = \frac{c_2}{w_2}$, there are infinitely many Nash equilibria as shown in Theorem 1.

The existence of multiple Nash equilibria in the game can be a challenge in practice. As Cachon and Netessine (2004) point out, it becomes a problem as no player knows which equilibrium will prevail and the outcome of the game cannot be predicted. They also state that there are some

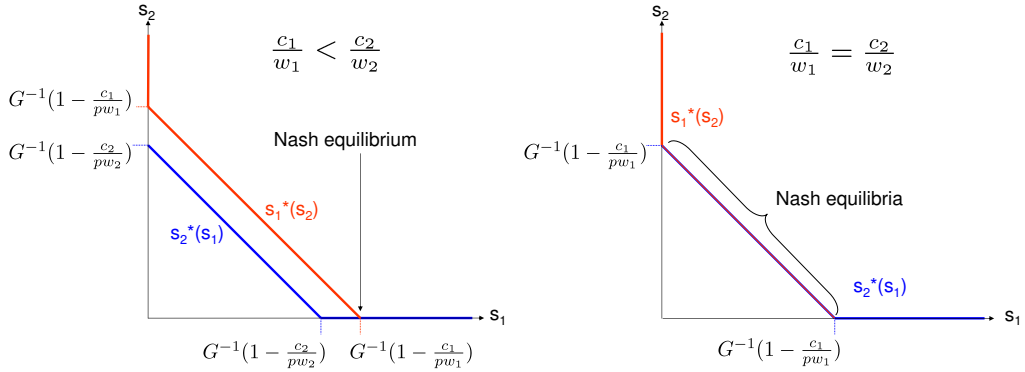


Figure 2: Nash equilibrium illustrated in the case with $n = 2$.

commonly used methods for handling multiple equilibrium solutions in a game, such as finding a Pareto optimal equilibrium or rules based on symmetry. As a result, in the presence of multiple Nash equilibria, the equilibrium that may seem preferable based on symmetry arguments is the symmetric equilibrium, i.e. the equilibrium where $\alpha_i = \frac{1}{m} \forall i = 1, \dots, m$. Indeed, according to Cachon and Netessine (2004), “a symmetric equilibrium is more focal than an asymmetric equilibrium” and a symmetric equilibrium may thus appear “more reasonable than others” (see also Mahajan and van Ryzin, 2001).

Cachon and Netessine (2004) also suggest that in the case of multiple equilibria, “an alternative method to rule out some equilibria is to focus only on the Pareto optimal equilibrium” as a “most preferred equilibrium by every player.” Sun et al. (2009), for example, also face a problem with multiple Nash equilibria, and focus on the Pareto optimal equilibrium which they prove to be unique. The next subsection explores Pareto optimality among the multiple Nash solutions.

2.4 Pareto Optimality

In this section, we assume that the Nash equilibrium is not unique, i.e. $m > 1$. Our goal is to determine whether one of the Nash equilibria is Pareto optimal. A Pareto optimal point is a solution such that reducing one of the players’ cost is impossible without increasing another player’s cost.

To find all the Pareto optimal points, it suffices to solve the following problem for all $\lambda_i \geq 0, \forall i$:

$$\begin{aligned}
& \min_{\mathbf{s}=(s_1, \dots, s_n) \geq 0} \Pi_P(\mathbf{s}) = \lambda_1 J_1(\mathbf{s}) + \dots + \lambda_{n-1} J_{n-1}(\mathbf{s}) + J_n(\mathbf{s}) \\
& = \min_{s_1, \dots, s_n \geq 0} \lambda_1 c_1 s_1 + \dots + \lambda_{n-1} c_{n-1} s_{n-1} + c_n s_n + (\lambda_1 w_1 + \dots + \lambda_{n-1} w_{n-1} + w_n) p E[U]. \quad (3)
\end{aligned}$$

Let $\lambda_n = 1$, $c'_i = \lambda_i c_i$, $\forall i$, and $p' = p \sum_{i=1}^n \lambda_i w_i$. Then problem (3) becomes

$$\begin{aligned} \min_{\mathbf{s}=(s_1, \dots, s_n) \geq 0} \Pi_P(\mathbf{s}) &= \sum_{i=1}^n c'_i s_i + p' E[U] \\ &= \sum_{i=1}^n c'_i s_i + p' \sum_{l=1}^L q^l \int_S^\infty (x - S) F^l(x) dx. \end{aligned}$$

Lemma 1. *We have $c'_i \leq p'$ $\forall i$.*

Proof. By Assumption 1, $c_i < p w_i$. Because $\lambda_i \geq 0$, $p > 0$ and $w_i > 0$, it follows that $c'_i = \lambda_i c_i \leq p \lambda_i w_i \leq p \sum_{j=1}^n \lambda_j w_j = p'$. \square

We want to find the Pareto optimal point that corresponds to a given set of values of λ_i , $i = 1, \dots, n$. Let $\mathcal{J} = \{i : c'_i = \min_j c'_j\}$ be the set of hospitals with minimum value of c'_i , and $j = |\mathcal{J}|$ the number of such hospitals. The following theorem provide a closed-form solution for Pareto optimal points.

Theorem 2. *If $j > 1$, there are infinitely many Pareto optimal points given by*

$$\hat{s}_i = \begin{cases} \gamma_i G^{-1}(1 - \frac{c'_i}{p'}) & i \in \mathcal{J} \\ 0 & i \notin \mathcal{J} \end{cases}$$

for any $\{\gamma_i, i \in \mathcal{J}\}$ such that $\gamma_i \in [0, 1]$ and $\sum_{i \in \mathcal{J}} \gamma_i = 1$.

Proof. Let $(\hat{s}_1, \dots, \hat{s}_n)$ an optimal solution, and let $\hat{S} = \sum_{i=1}^n \hat{s}_i$ and $\hat{S}_{-i} = \hat{S} - \hat{s}_i$. The partial derivatives of $\Pi_P(s_1, \dots, s_n)$ are

$$\begin{aligned} \frac{\partial \Pi_P}{\partial s_i}(s_1, \dots, s_n) &= c'_i - p' \sum_{l=1}^L q^l \int_S^\infty f^l(x) dx \\ &= c'_i - p' \sum_{l=1}^L q^l (1 - F^l(S)) \\ &= c'_i - p' + p' \sum_{l=1}^L q^l F^l(S) \\ \frac{\partial \Pi_P}{\partial s_i \partial s_j}(s_1, \dots, s_n) &= p' \sum_{l=1}^L q^l f^l(S). \end{aligned}$$

It is clear that the Hessian matrix of Π_P is positive semi-definite for all (s_1, \dots, s_n) , so Π_P is convex. We note that a convex function of a variable attains its minimum on \mathbb{R}^+ at zero if its derivative at zero is positive, and at a point where its derivative equals zero otherwise. As a result, we have:

- if $c'_i - p' + p'G(\hat{S}_{-i}) > 0$, then $\hat{s}_i = 0$;
- if $c'_i - p' + p'G(\hat{S}_{-i}) \leq 0$, then \hat{s}_i satisfies $c'_i - p' + p'G(\hat{S}) = 0$.

In other words, $\hat{s}_i = G^{-1}(1 - \frac{c'_i}{p'}) - \hat{S}_{-i}$ if this quantity is non-negative, and $\hat{s}_i = 0$ otherwise. The result follows using a reasoning similar to the proof of Theorem 1. \square

The following corollary can be proved similarly (and the proof is thus omitted).

Corollary 3. *If $j = 1$, there is a unique Pareto optimal point given by*

$$\hat{s}_i = \begin{cases} G^{-1}(1 - \frac{c'_i}{p'}) & \text{if } i \in \mathcal{J} \\ 0 & \text{if } i \notin \mathcal{J}. \end{cases}$$

The closed-form expression for the total stockpile S_P at a Pareto optimal point corresponding to a given set of values of λ_i , $i = 1, \dots, n$ follows directly from the results above.

Corollary 4. *The system-wide total stockpile level S_P at a Pareto optimal point corresponding to a given set of values of λ_i , $i = 1, \dots, n$ is given by*

$$S_P = \sum_{i=1}^n \hat{s}_i = G^{-1}(1 - \frac{c'_{i_1}}{p'}) \quad \forall j, \gamma_i, i \in \mathcal{J},$$

where c'_{i_1} is the minimum of all the coefficients c'_i , i.e., $c'_{i_1} = \min_i c'_i$.

In particular, the total stockpile at a Pareto optimal point is the same at all Pareto optimal points corresponding to the same λ_i , $i = 1, \dots, n$, even though the individual stockpiles at distinct solutions are not the same. Note however that the total stockpile depends on λ_i , $i = 1, \dots, n$.

2.4.1 Particular case: A two-hospital game

Let us consider a two-hospital setting. As explained above, the study of Pareto optimal points is of particular importance in the case of multiple Nash equilibria, so we assume that the Nash

equilibrium is not unique, i.e. $\frac{c_1}{w_1} = \frac{c_2}{w_2}$.

Proposition 3. *There is no Nash equilibrium that is Pareto optimal in a two-hospital game with multiple Nash equilibria.*

Proof. First we recall that under the assumption $\frac{c_1}{w_1} = \frac{c_2}{w_2}$, the Nash equilibria correspond to any arbitrary split of $S_d = G^{-1}(1 - \frac{c_1}{pw_1}) = G^{-1}(1 - \frac{c_2}{pw_2})$ among \bar{s}_1 and \bar{s}_2 . We consider three cases.

Case 1: $0 \leq \lambda < \frac{c_2}{c_1}$, i.e. $c'_1 < c'_2$ so there is a unique Pareto optimal point associated with λ : $\hat{s}_1 = G^{-1}(1 - \frac{c'_1}{p'})$, $\hat{s}_2 = 0$. This point is a Nash equilibrium iff $G^{-1}(1 - \frac{c'_1}{p'}) = G^{-1}(1 - \frac{c_1}{pw_1})$, or equivalently

$$\frac{\lambda c_1}{p(\lambda w_1 + w_2)} = \frac{c_1}{pw_1},$$

which is impossible as $w_2 \neq 0$.

Case 2: $\lambda > \frac{c_2}{c_1}$ (in particular $\lambda \neq 0$), i.e. $c'_1 > c'_2$ so there is a unique Pareto optimal point associated with λ : $\hat{s}_1 = 0$, $\hat{s}_2 = G^{-1}(1 - \frac{c'_2}{p'})$. This point is a Nash equilibrium iff $G^{-1}(1 - \frac{c'_2}{p'}) = G^{-1}(1 - \frac{c_2}{pw_2})$, or equivalently

$$\frac{c_2}{p(\lambda w_1 + w_2)} = \frac{c_2}{pw_2},$$

which is impossible as $\lambda w_1 \neq 0$.

Case 3: $\lambda = \frac{c_2}{c_1}$, i.e. $c'_1 = c'_2$ so there is an infinite number of Pareto optimal points associated with λ , corresponding to any arbitrary split of $S_P = G^{-1}(1 - \frac{c'_1}{p'}) = G^{-1}(1 - \frac{c'_2}{p'})$ among \hat{s}_1 and \hat{s}_2 . One of these points is a Nash equilibrium iff they all are Nash equilibria, which is the case iff $G^{-1}(1 - \frac{c'_1}{p'}) = G^{-1}(1 - \frac{c_1}{pw_1})$, which is impossible as $w_2 \neq 0$. \square

The following proposition compares the total stockpile at a Pareto optimal point and at a Nash equilibrium.

Proposition 4. *The total stockpile level at any Pareto optimal point is greater than at a Nash equilibrium in a two-hospital game with multiple Nash equilibria.*

Proof. If $\lambda < \frac{c_2}{c_1}$, then since G^{-1} is increasing and $w_2 > 0$,

$$S_P = G^{-1}(1 - \frac{c'_1}{p'}) = G^{-1}(1 - \frac{\lambda c_1}{p(\lambda w_1 + w_2)}) > G^{-1}(1 - \frac{c_1}{pw_1}).$$

If $\lambda > \frac{c_2}{c_1}$, then since G^{-1} is increasing and $w_1 > 0$,

$$S_P = G^{-1}\left(1 - \frac{c_2'}{p'}\right) = G^{-1}\left(1 - \frac{c_2}{p(\lambda w_1 + w_2)}\right) > G^{-1}\left(1 - \frac{c_2}{p w_2}\right).$$

If $\lambda = \frac{c_2}{c_1}$, then since $\frac{c_1}{w_1} = \frac{c_2}{w_2}$,

$$S_P = G^{-1}\left(1 - \frac{c_1'}{p'}\right) = G^{-1}\left(1 - \frac{\lambda c_1}{p(\lambda w_1 + w_2)}\right) = G^{-1}\left(1 - \frac{1}{p\left(\frac{w_1}{c_1} + \frac{w_2}{c_2}\right)}\right) = G^{-1}\left(1 - \frac{c_1}{2p w_1}\right) > G^{-1}\left(1 - \frac{c_1}{p w_1}\right).$$

□

Figure 3 illustrate the set of Pareto points in the case of two hospitals with multiple Nash equilibria.




Figure 3: Nash equilibria and Pareto optimal points illustrated in the case with $n = 2$.

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. We note that the centralized setting does not allow any risk pooling because the stochasticity of the demand is introduced at the aggregate level via the total demand for the entire group of hospitals, and not at the individual hospitals level. The goal of the central planner is then to decide the stockpiles that minimize the expected total cost for the n hospitals:

$$\min_{\mathbf{s}=(s_1,\dots,s_n)\geq 0} \Pi_c(\mathbf{s}) = \sum_{i=1}^n J_i(s_i, \mathbf{s}_{-i}) = \sum_{i=1}^n c_i s_i + p \sum_{l=1}^L q^l \int_S^\infty (x - S) f^l(x) dx. \quad (4)$$

It follows from Assumption 1 that $c_i < p$, $i = 1, \dots, n$.

Let $\mathcal{I} = \{i : c_i = \min_j c_j\}$ be the set of hospitals with minimum unit stockpiling cost, and $k = |\mathcal{I}|$ the number of hospitals with minimum unit cost. The following theorem provide a closed-form solution for the centralized problem (4).

Theorem 3. *If $k > 1$, there are infinitely many centralized solutions given by*

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

for any $\{\beta_i, i \in \mathcal{I}\}$ such that $\beta_i \in [0, 1]$ and $\sum_{i \in \mathcal{I}} \beta_i = 1$.

The proof is similar to the proof of Theorem 2 and is thus omitted.

The following corollary can be proved similarly to Corollary 3.

Corollary 5. *If $k = 1$, there is a unique centralized solution given by*

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

The interpretation behind these two results is simple. 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 the above results.

Corollary 6. *The system-wide total stockpile level at any centralized solution, S_c , is given by*

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

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 particular, the total stockpile at a centralized solution is the same at all centralized solutions, even though the individual stockpiles at distinct solutions are not the same.

The centralized solutions (as well as Nash equilibrium solutions) are illustrated in Figure 4 in the case of two hospitals in several different cases.

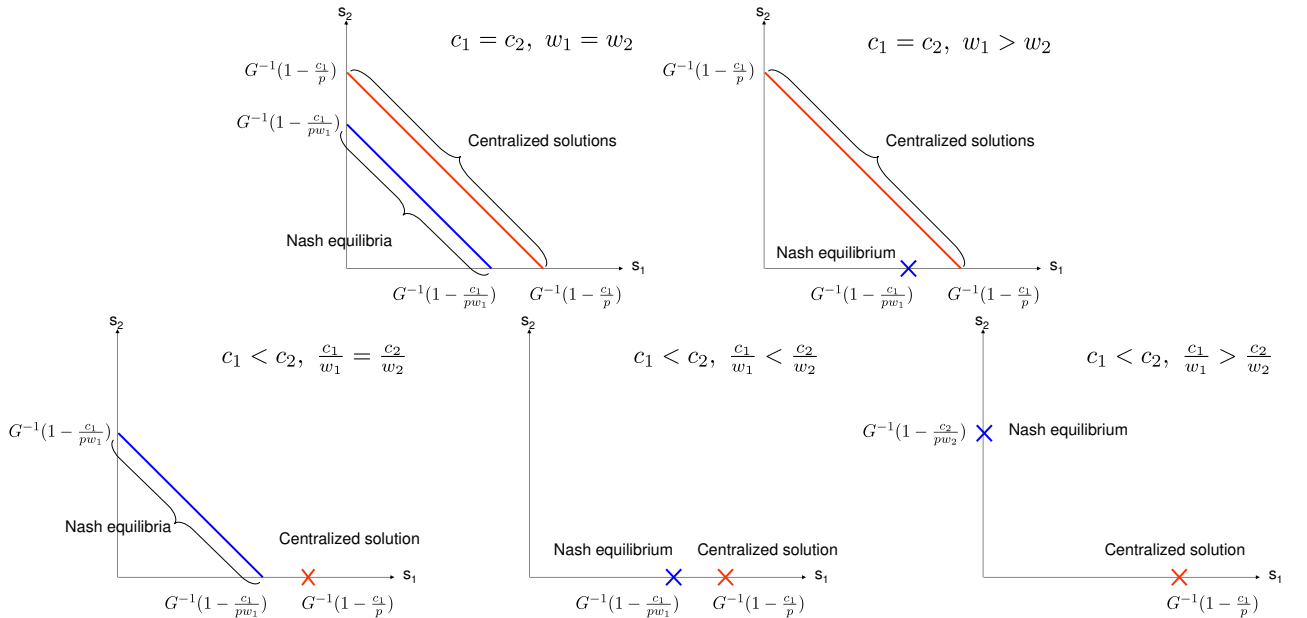


Figure 4: Centralized solution and Nash equilibrium illustrated in the case with $n = 2$.

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 β_i , $i \in \mathcal{I}$ 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

Let us 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 that the equilibrium may be very inefficient in terms of cost, when its performance is compared with what could be achieved should a coordinated decision be made.

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

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

From Corollaries 2 and 6, we have

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

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.

The worst loss of efficiency in quantity and the loss of efficiency in cost are illustrated in Figure 5 in the case of two hospitals, $L = 3$ scenarios ($q^1 = 0.2$, $q^2 = 0.5$, $q^3 = 0.3$) where the demand distribution follows a Normal distribution with mean of 100, 300 and 500 and standard deviation of 30, 70 and 120 in each of the three scenarios respectively.

4.2 Aggregate amount of stockpiles

We now focus on the ratio σ_d that compares the aggregate amount of stockpile in the centralized and decentralized settings.

Proposition 5. *The total stockpile at a Nash equilibrium is lower than the total stockpile at a centralized solution: $S_d < S_c$ (i.e. $\sigma_d > 1$).*

Proof. The result follows from the fact that $c_{i_0} \leq c_1 < \frac{c_1}{w_1}$ and that G^{-1} is strictly increasing. \square

This result is illustrated in Figure 5 where we observe that σ_d is always greater than 1. It implies that if a region decides to centralize its stockpiling decisions, it would always result in a larger total stockpile than if each hospital makes its own decision by minimizing its expected cost with shared resources. Therefore, centralizing the decision would not only reduce the overall costs, but it would also yield a larger overall stockpile.

4.3 Sensitivity to the penalty coefficient

It is clear from the closed-form expressions obtained that the total stockpile quantities in both settings are non decreasing with the parameter p as a higher penalty value gives incentives to stockpile more. We now investigate how the efficiency ratios vary with this penalty coefficient.

Lemma 2. *As p becomes infinite, σ_d tends to 1.*

Proof. The result is clear if b is finite as $\lim_{p \rightarrow \infty} S_d = \lim_{p \rightarrow \infty} S_c = G^{-1}(1) = b$. We now assume $b = \infty$, i.e. $\lim_{x \rightarrow 1} G^{-1}(x) = +\infty$. We have

$$1 - \frac{c_{i_0}}{p} = 1 - \frac{c_1}{pw_1} + \frac{1}{p} \left(\frac{c_1}{w_1} - c_{i_0} \right).$$

As $p \rightarrow \infty$, the second term above becomes arbitrarily small. Therefore, for a large p ,

$$G^{-1}\left(1 - \frac{c_{i_0}}{p}\right) \simeq G^{-1}\left(1 - \frac{c_1}{pw_1}\right) + \frac{1}{p} \left(\frac{c_1}{w_1} - c_{i_0} \right) (G^{-1})' \left(1 - \frac{c_1}{pw_1}\right) = G^{-1}\left(1 - \frac{c_1}{pw_1}\right) + \frac{1}{p} \frac{\frac{c_1}{w_1} - c_{i_0}}{\sum_{l=1}^L q^l f^l \left(G^{-1}\left(1 - \frac{c_1}{pw_1}\right)\right)}.$$

Therefore, for a large p ,

$$\sigma_d \simeq 1 + \frac{1}{p} \frac{\frac{c_1}{w_1} - c_{i_0}}{G^{-1}\left(1 - \frac{c_1}{pw_1}\right) \sum_{l=1}^L q^l f^l \left(G^{-1}\left(1 - \frac{c_1}{pw_1}\right)\right)}.$$

Now, $\lim_{p \rightarrow \infty} \frac{1}{p} = 0$, $\lim_{p \rightarrow \infty} \frac{\frac{c_1}{w_1} - c_{i_0}}{G^{-1}\left(1 - \frac{c_1}{pw_1}\right)} = 0$ and $\lim_{p \rightarrow \infty} \frac{1}{\sum_{l=1}^L q^l f^l \left(G^{-1}\left(1 - \frac{c_1}{pw_1}\right)\right)} < \infty$, therefore $\lim_{p \rightarrow \infty} \sigma_d = 1$. \square

Figure 5 illustrates this result. This implies that while the centralized total stockpile is always greater than the decentralized total stockpile, for a very large value of the penalty parameter, the stockpile amount at the decentralized equilibrium becomes close to the centralized stockpile amount. As a result, in a setting where each hospital makes its own decisions and sharing or demand redistribution apply, a public policy that renders it very costly to have shortages would not only lead to a higher stockpile level, but also would make the system as well prepared as if the decisions had been made in a coordinated fashion.

It is also worth pointing out that the worst loss of efficiency in cost for the decentralized setting is not monotonic with p , as shown in Figure 5. Indeed, an increase in the penalty value does not always decrease the worst loss of efficiency in cost. The rationale of this observation is that if $c^{max} \neq c_{i_0}$, the hospital that stockpiles at the centralized solution is not the one that stockpiles at the Nash solution. In this case, the purchasing cost at the “worst” Nash equilibrium is given by $c^{max} S_d$, and the purchasing cost at the centralized solution is given by $c_{i_0} S_c$. We have $\lim_{p \rightarrow \infty} S_d = \lim_{p \rightarrow \infty} S_c = G^{-1}(1) = b$, but because $c^{max} \neq c_{i_0}$, $\lim_{p \rightarrow \infty} c^{max} S_d \neq \lim_{p \rightarrow \infty} c_{i_0} S_c$ and therefore the expected stockpile costs do not converge to the same value as the unit penalty

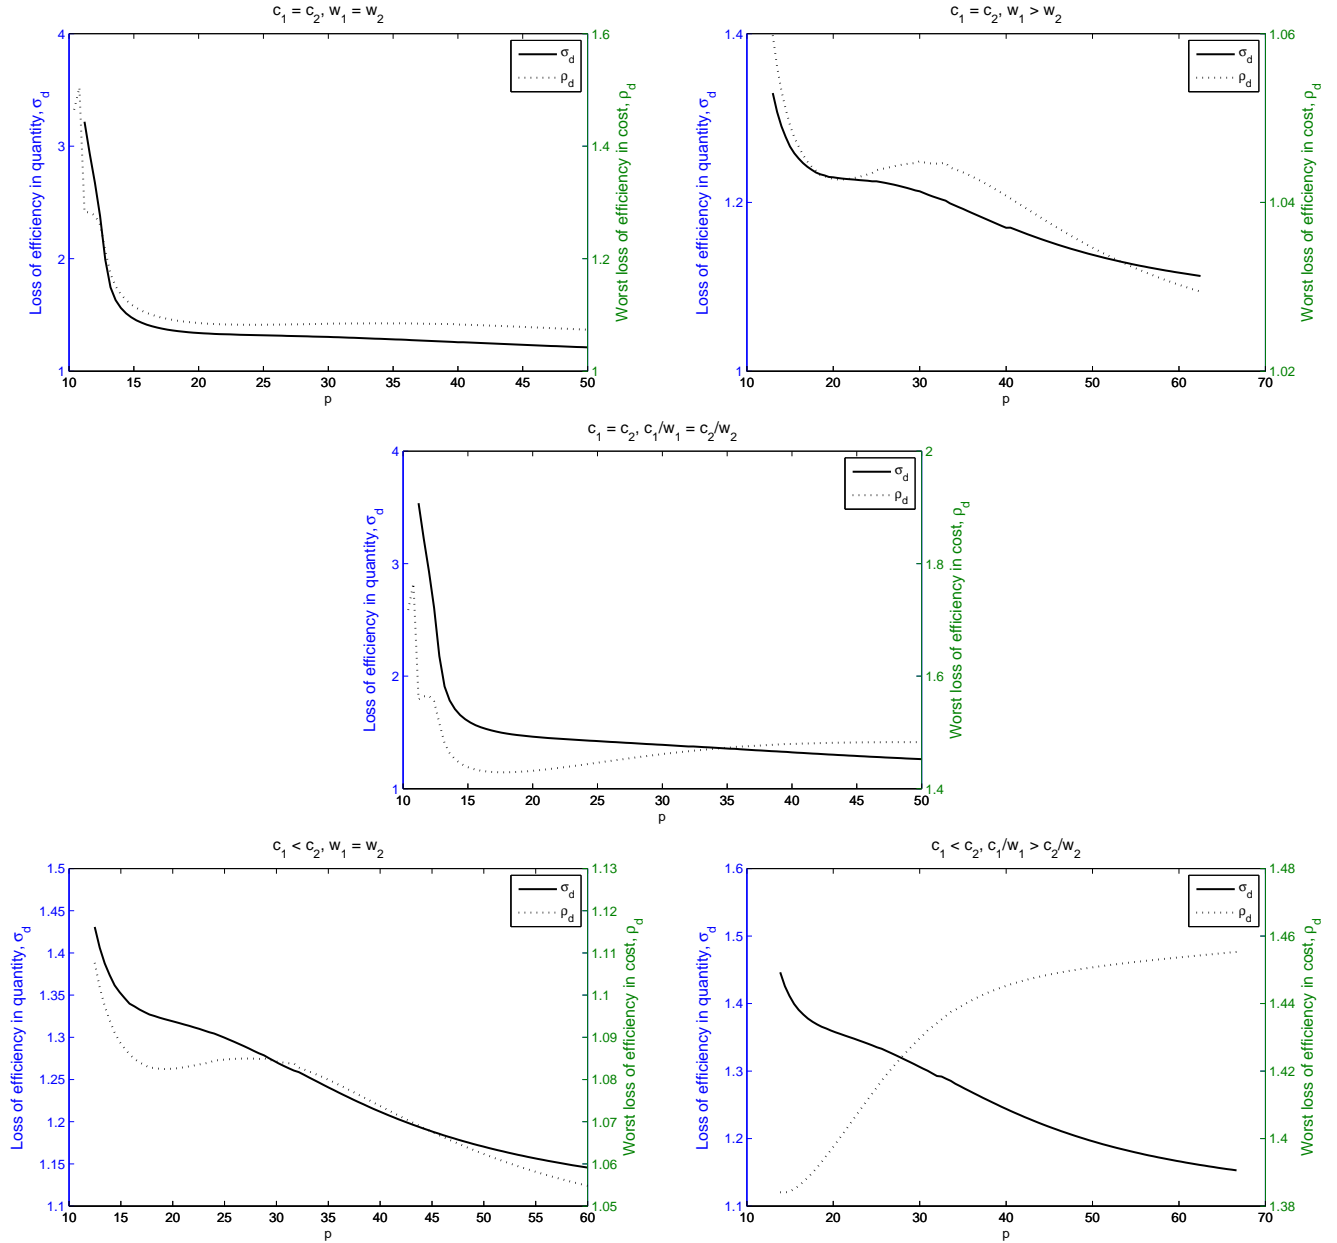


Figure 5: The loss of efficiency in quantity, σ_d , and the worst loss of efficiency in cost, ρ_d , are illustrated here for two hospitals ($n = 2$). Parameters (c_1, c_2, w_1, w_2) are given respectively: $(5, 5, 0.5, 0.5)$, $(5, 5, 0.6, 0.4)$, $(4, 6, 0.4, 0.6)$, $(4, 6, 0.5, 0.5)$ and $(4, 6, 0.3, 0.7)$.

cost approaches infinity.

Although in Figure 5, σ_d appears monotonically decreasing in p , another example depicted in the left-hand-side graph in Figure 6 shows that this is not generally the case. Similarly, the worst loss of efficiency in cost, ρ_d , is not in general monotonically decreasing in p as shown in the right-hand-side plot in Figure 6.




Figure 6: A two-hospital example showing that the ratio of loss of efficiency in cost and quantity is not necessarily monotonically decreasing in the penalty cost p . Parameters used: $L = 3$, demand distributions with means = [100 300 500] and standard deviations [30 70 120], $c = [4 6]$, $w = [0.4 0.6]$, $q = [0.8 0.1 0.1]$.

4.4 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 (i.e.

economies of scale). A strategically located warehouse and efficiently managed stockpile can also help reduce the supply unit cost.

4.5 Sensitivity to the weight parameters

The effect of the weight parameter w_i of hospital i on its stockpile level is quite straightforward. First of all, we know that the weights w_i do not affect S_c . Also, clearly, as w_1 increases, $S_d = G^{-1}(1 - \frac{c_1}{pw_1})$ increases. Intuitively, the centralized aggregate stockpile is independent of the weights of each hospital as only the overall penalty matters. At the decentralized solution, the aggregate stockpile goes up when the hospital(s) that actually stockpile is(are) responsible for a larger share of the overall penalty.

4.6 Insights

Based on the characteristics of the centralized and Nash solutions, some managerial insights can be drawn for disaster preparedness purposes. First of all, in a coordinated setting, a higher aggregate stockpile quantity will result than in a decentralized setting. This implies that a centralized approach leads to a better prepared community than what may result in the decentralized setting in which each hospital makes its own decision considering supply sharing.

As Lemma 2 shows, the loss of efficiency in quantity becomes arbitrarily close to one for large values of the unit penalty cost. This means that if the perceived unit penalty cost due to supply shortage can be made very large, e.g. because of the serious threat of a disaster with high mortality rate, or because the government is imposing a monetary penalty for supply shortages, σ_d would become close to one, meaning that the aggregate stockpile amount would be close to that of the centralized setting.

From the analysis of the sensitivity to unit supply cost c_i we can conclude that a discounted unit cost leads to a higher total stockpile. Such a discount can be achieved by government provided subsidy for disaster planning purposes, or by hospitals forming a large purchasing power in order to negotiate a better pricing with their suppliers.

It is worth noting that in this problem, it may be reasonable to assume that all hospitals in the group have similar unit costs since in the healthcare industry, there are few suppliers/distributors that carry a given item in a region, and hospitals typically purchase medical supplies in groups

(e.g. by affiliation) but not individually. Thus, in the decentralized setting, the hospital with the largest stake in overall shortage penalty (i.e. largest w_i) will stockpile all the group needs (since it has the smallest ratio of $r_i = c_i/w_i$).

In summary, when the decision-making for hospital stockpiling in preparation for disasters is done in a centralized manner, the aggregate cost is lower and the aggregate stockpile quantity is higher, hence the community is better prepared for disaster with a lower cost. The optimal centralized solution has only the hospital(s) with lowest unit cost purchase supplies, but in practice this stockpiling cost would be distributed among all participating hospitals. In contrast, when the decision-making is decentralized, only the hospital(s) with the lowest ratio of cost over weight will stockpile for the entire community, yielding higher overall costs and a lower aggregate stockpile. This incurs a large sunk cost for the hospital(s) that do stockpile, which is likely recurrent every few years due to the need for stock rotation while all other hospitals in the group do not invest at all prior to the onset of a disaster. Therefore, from the perspective that disaster preparedness should be a community-wide, shared effort aiming at best serving the need of a population when a disaster occurs, the decentralized setting presents serious drawbacks. There exist several ways for the government to remedy the situation such as by providing subsidies to help lower the cost of stockpiling. Another possibility is imposing extra shortage penalties as a means to give more incentives for hospitals to stockpile larger quantities.

5 Conclusion

Sufficient hospital stockpile of medical supplies in preparation for a disaster is a task demanding immediate attention and action. We present this problem as 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 a mutual aid agreement for sharing resources. This model 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. In particular, we found that

increasing the shortage penalty or decreasing the procurement cost would improve the efficiency of the equilibrium, in the sense that the equilibrium stockpile would be larger and closer to the stockpile of the system optimum. This could be achieved via a government intervention such as imposing a tax on supplies shortage during a disaster or subsidizing the purchase of supplies to encourage adequate planning. The efficiency of the equilibrium may also be improved via contracts among participating hospitals and possibly government.

Although our modeling approach relies on a number of assumptions that may not hold in a strict sense in reality, we feel that the solutions we obtain reflect the reality that our partner hospital has observed in the recent past while preparing for an influenza pandemic: this large hospital has been stockpiling supplies while some other smaller hospitals in the area had decided not to do so. Therefore, we feel that the assumptions and model have been set up appropriately to reflect the main characteristics of the problem, while remaining tractable enough to allow for an analytical approach.

One extension of this work is evaluating the performance of Nash stockpile levels and the centralized solutions when real demand distributions deviate from the assumed ones. In the current model, hospitals are assumed to have full knowledge of the probability of each disaster scenario and its demand distribution. However, in reality there is little scientific evidence available on how to estimate the severity a specific disaster and its probability of occurrence. Future research will involve relaxing the assumption that the exact probability distribution is known.

Another extension is setting up the problem as a cooperative game. In a cooperative game, hospitals form coalitions and determine the best binding agreement that achieves the optimum of its objective. In such a game, the decisions of each hospital can be the fraction of excess supply it shares with others when needed. This paper takes a non-cooperative game approach to represent our observations of the way hospitals make decisions – individually so as to maximize their own utility. However, considering a cooperative approach could reveal interesting insights and suggest a potential alternative to how decisions are currently being made.

Our stockpile game is a simultaneous game. If hospitals make stockpile decisions sequentially, we expect the resulting stockpiles of each hospital to be very different. This extension can shed light especially on the leader-follower behavior among hospitals of interest.

Finally, another possible direction of extension is to design implementable contracts among

hospitals and possibly involving government in a transfer payment scheme such that each hospital's objective becomes aligned with the system's objective. It then improves the performance of the equilibrium and achieves coordination of the system, thereby avoiding inefficiencies due to the decentralization of the decision-making.

References

- Cachon, G. and Netessine, S. (2004) Game theory in supply chain analysis. In *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, Simchi-Levi, D., Wu, S. D., and Shen, Z.-J. M., eds. Kluwer, Norwell, MA, pp. 13–65.
- Chan-Yeung, M. (2004) Severe acute respiratory syndrome (SARS) and healthcare workers. *International Journal of Occupational and Environmental Health*, **10**(4), 421–427.
- Chick, S. E., Mamani, H., and Simchi-Levi, D. (2008) Supply chain coordination and influenza vaccination. *Operations Research*, **56**(6), 1493–1506.
- Havlak, R., Gorman, S. E., and Adams, S. A. (2002) Challenges associated with creating a pharmaceutical stockpile to respond to a terrorist event. *Clinical Microbiology and Infectious Diseases*, **8**, 529–533.
- Hung, L. (2003) The SARS epidemic in Hong Kong: what lessons have we learned? *Journal of the Royal Society of Medicine (JRSM)*, **96**(8), 374.
- Jacobson, S. H., Sewell, E. C., Proano, R. A., and Jokela, J. A. (2006) Stockpile levels for pediatric vaccines: how much is enough? *Vaccine*, **24**, 3530–3537.
- Julian Fantino (2003) SARS Outbreak: The Response of the Toronto Police Service. http://policechiefmagazine.org/magazine/index.cfm?fuseaction=display_arch&article_id=565&issue_id=42005. Accessed: May 13, 2010.
- Krishnan, K. S. and Rao, V. R. K. (1965) Inventory control in N warehouses. *Journal of Industrial Engineering*, **16**, 212–215.

- Lee, V. J., Phua, K. H., Chen, M. I., Chow, A., Ma, S., Goh, K. T., and Leo, Y. S. (2006) Economics of neuraminidase inhibitor stockpiling for pandemic influenza, Singapore. *Emerging Infectious Diseases*, **12**(1), 95–102.
- Mahajan, S. and van Ryzin, G. (2001) Inventory competition under dynamic consumer choice. *Operations Research*, **49**(5), 646–657.
- Medema, J. K., Zoellner, Y. F., Ryan, J., and Palache, A. M. (2004) Modeling pandemic preparedness scenarios: health economic implications of enhanced pandemic vaccine supply. *Virus Research*, **103**, 9–15.
- Rudi, N., Kapur, S., and Pyke, D. F. (2001) A two-location inventory model with transshipment and local decision making. *Management Science*, **47**(12), 1668–1680.
- Sun, P., Yang, L., and de Véricourt, F. (2009) Selfish drug allocation for containing an international influenza pandemic at the onset. *Operations Research*, **57**(6), 1320–1332.
- Tagaras, G. (1989) Effects of pooling on the optimization and service levels of two-location inventory systems. *IIE Transactions*, **21**, 250–257.
- The Centers for Disease Control and Prevention (CDC) (2008) Health, United States, 2008. <http://www.cdc.gov/nchs/hus.htm>. Accessed: December 16, 2009.
- The Centers for Disease Control and Prevention (CDC) (2009) 2009-2010 Influenza Season Week 40 ending October 10, 2009. <http://www.cdc.gov/flu/weekly/weeklyarchives2009-2010/weekly40.htm>. Accessed: December 16, 2009.
- Toner, E. and Waldhorn, R. (2006) What hospitals should do to prepare for an influenza pandemic. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, **4**(4), 397–402.
- Wang, S., de Véricourt, F., and Sun, P. (2009) Decentralized resource allocation to control and epidemic: A game theoretic approach. *Mathematical Biosciences*, **222**, 1–12.
- Zhao, H., Deshpande, V., and Ryan, J. K. (2005) Inventory sharing and rationing in decentralized dealer networks. *Management Science*, **51**(4), 531–547.