Journal Pre-proof

Cybersecurity Investments in the Supply Chain: Coordination and a Strategic Attacker

Jay Simon, Ayman Omar

PII: S0377-2217(19)30757-X
DOI: https://doi.org/10.1016/j.ejor.2019.09.017
Reference: EOR 16045

To appear in: *European Journal of Operational Research*

Received date: 24 January 2019
Accepted date: 10 September 2019

Please cite this article as: Jay Simon, Ayman Omar, Cybersecurity Investments in the Supply Chain: Coordination and a Strategic Attacker, *European Journal of Operational Research* (2019), doi: https://doi.org/10.1016/j.ejor.2019.09.017

Highlights

- Lack of supply chain coordination leads to underinvestment in cybersecurity
- Underinvestment is partly mitigated by the presence of a strategic attacker
- Higher indirect damages from an attack are associated with more severe underinvestment
- Overinvestment can occur if indirect damages from an attack are relatively minor

Cybersecurity Investments in the Supply Chain: Coordination and a Strategic Attacker

Jay Simon
American University
Kogod School of Business
4400 Massachusetts Ave, NW
Washington, DC 20016
Corresponding author: jaysimon@american.edu

Ayman Omar
American University
Kogod School of Business
4400 Massachusetts Ave, NW
Washington, DC 20016

Abstract

Cybersecurity poses a difficult challenge to supply chains, as a firm may be affected by an attack on another firm in the supply chain. For example, a retailer's consumer data might be compromised via an attack on a supplier. In general, individual nodes in a supply chain bear the entire cost of their own cybersecurity investments, but some of the benefits of the investments may be enjoyed by the other nodes as well. We analyze the differences between coordinated and uncoordinated cybersecurity investments, as well as the differences resulting from a strategic and a non-strategic attacker. We find that lack of coordination leads to underinvestment with a non-strategic attacker, but that this is somewhat counterbalanced by an attacker being strategic. Lack of coordination may lead to either underinvestment or overinvestment with a strategic attacker, depending on how large the indirect damages from attacks are relative to the direct damages; overinvestment is more likely if indirect damages are relatively minor. A numerical example is provided to illustrate the impacts of and relationships between coordinated investments and a strategic attacker.

Keywords: supply chain management; cybersecurity; supply chain coordination; attacker-defender model; interdependent security

1 Introduction

Today's global supply chains are exposed to diverse risks, any of which can temporarily disrupt a firm's operations (Rao and Goldsby, 2009). One of the emerging risks in supply chains deals with cybersecurity concerns, which must be addressed in any broader supply chain risk management strategy (e.g. the portfolio approach of Kumar and Park (2019)). Cybersecurity is "the body of combined technologies, processes, and practices that are put in place to protect data and networks from attacks, damage, or unauthorized access" (Boyson, 2014). Cybercrime has a significant negative impact on the world economy with annual costs of about 445 billion US dollars (McAfee Center for Strategic and International Studies, 2014). The relevance of cyber risks and their impacts have increased substantially during the last several years. This was also emphasized in Allianz Global Corporate and Specialty (2017), where cyber risks placed in the top three business risks, versus being ranked 15th in 2013. The frequency and impact of cyberattacks have increased in the United States at an annual rate of 25% over the last three years. This same trend has also been taking place in Europe, where several countries have experienced a growth in successful breaches and significant damages (Smith, 2016; Mahwah, 2017). Such attacks have targeted organizations in both the public and the private sector. One of the most recent successful breaches of major corporations involved Marriott, where identities of 500 million customers were exposed, an attack that was discovered four years after it took place.

A cyberattack on a supplier or third party logistics can cause significant delays in production schedules and supply, and can result in delays of shipments and deliveries across supply chains (Williams, 2014). Cyber breaches through third party providers can also have an impact on the overall financial performance of an organization and its shareholders value (Modi et al., 2015). Such attacks do not have to be directed at the organization itself, but can originate anywhere in the supply chain. There are many cases where companies were affected due to a breach of a third party service provider. Organizations such as Target, Home Depot, Fiat Chrysler, T-Mobile USA, the IRS, CVS, Costco, Sam's Club, Boston Medical Center, and others have suffered cyberattacks because their third party providers were compromised. It is concerning that third party breaches have been steadily increasing; a report by AIG (2013) estimates that more than 60% of breaches take place through a third party provider. Supply chains' reliance on information technology has increased significantly in recent years and will continue to do so in the future. Emerging technologies and trends, such as 3D-printing, e-commerce, Industry 4.0, internet of things, physical internet, and smart homes have increased the points of vulnerability resulting in a higher likelihood for cyberattacks.

Damages resulting from a successful cyberattack can take many different forms. There can be a major impact on an organization's brand and financial performance (Modi et al., 2015). Cyberattacks can lead to significant damages due to lawsuits, legal fees, and the loss of customer goodwill and trust. For example, when Target suffered a breach in customer data in 2013, their total

financial damages were over $200 million (Ramakrishnan and Bose, 2017). Such breaches can be carried out in a matter of minutes or hours, but take months or years to detect and contain (Smith, 2016). Attacks can cause major disruptions in logistics, production and operations, as well as loss of data (Yang and Wei, 2013). Interruption in operations can lead to additional costs to several companies in the supply chain to return to the original state of operations (Davis, 2015). Attacks launched at suppliers can disrupt operations leading to a cascading negative impact on the entire supply chain. In addition to delays, this can lead to a lack of available inventory, as well as costs from expedited shipping needs to counter the disruption. An attack on a supply chain node can also result in the loss of intellectual property for a company, or a reduction in service levels to the end customers (Khan and Estay, 2015).

In the model analyzed in this paper, we deliberately use generic damage parameters to avoid limiting its applicability to specific types of attacks or organizations. A particular application of the model might include any of the forms of damage from a cyberattack given above.

Recent reports show that supply chain managers lack awareness in managing cyber risks (Gaudenzi and Siciliano, 2017). An important aspect of managing this risk is allocating the appropriate investments. Managing a company's cyber risks and threats can no longer be viewed in isolation, rather this should be addressed as a problem for the entire supply chain.

Existing research addresses some parts of managing cyber risks, yet a significant gap in the literature remains. The literature is rich with studies that examine risk types and risk mitigation strategies, and in some instances look at investment levels. However, relatively little work has been done to explore the appropriate investment levels beyond two organizations in the supply chain. More importantly, prior work tends not to differentiate investment allocations between the cases of a strategic versus non-strategic attacker. A strategic attacker directs a breach attempt at a deliberately-chosen organization, and is influenced by cybersecurity investments. A non-strategic attacker's choice of target in the supply chain is not affected by these investments. This distinction has implications to the risk profile of the supply chain, and accordingly to the optimal investment levels to manage the risk.

The objective of this research is to address the existing gap in the literature by examining optimal cybersecurity investments in supply chains using two dimensions. The dimensions are the type of attacker (strategic versus non-strategic) and the level of coordination of investments across nodes in the supply chain (coordinated versus non-coordinated supply chains).

We first obtain some general results that are consistent with prior literature. If nodes are uncoordinated, they will systematically underinvest in cybersecurity, because they ignore the positive externalities of their own investments. A simple coordination mechanism can resolve this problem and induce each node to invest at the supply chain optimal level.

In addition, if nodes are permitted to invest in cybersecurity for *other* nodes, we find that it can be optimal for larger nodes to invest in cybersecurity for smaller nodes unilaterally, even in the absence of coordination. This phe-

nomenon arises in real-world supply chains. For example, a successful attack on Aramco, the giant Saudi oil company, halted operations and led to major changes in how the company dealt with its third party providers' cybersecurity readiness (Woods and Bochman, 2018). Aramco began hiring experts to assess the readiness of its contractors and, if needed, address some of the IT security concerns.

We then explore the possibility of a strategic attacker, i.e. where nodes on which an attack would lead to a higher expected total damage are more likely to be attacked. This introduces another externality to cybersecurity investments: in addition to reducing indirect damages to other nodes from an attack on a given node, they also "push" some attack probability from that given node onto other nodes. These two externalities counterbalance one another. Depending on how large the spillover effects of an attack are relative to the direct impact on the node attacked, it is possible to observe either underinvestment or overinvestment when the attacker is strategic.

The remainder of the paper proceeds as follows. In Section 2, we review prior work related to our analysis. In Section 3, we develop and analyze a model using a non-strategic attacker, then introduce a strategic attacker, and explore the supply chain implications of each type. We present a numerical example in Section 4. Section 5 concludes the paper.

2 Literature Review

In addition to the work mentioned in the Introduction, there is substantial prior research relevant to this paper, both on supply chain security and on attacker-defender games.

Supply chains are interdependent nodes that work together to provide a good or service. Hausken and Levitin (2012) provide a classification of articles on systems defense and attack models; one of their categories of system structure is "interdependent systems," in which "an impact on one element gets transferred further to one or several other elements due to linkages" (p. 356), which is a key characteristic of the current paper. Accordingly, the importance and criticality of identifying the effective choices of investments becomes even higher in such systems. The concepts of interconnectedness and interdependence have been emphasized by Chopra and Khanna (2015), who analyze an economic input-output model of infrastructure sectors in the United States based on empirical data. They find that the overall system is relatively robust to random failures, but vulnerable to targeted disruptions. Wu et al. (2016) analyze system vulnerabilities in the context of terrorist attacks on infrastructure networks, considering both physical and geographical interdependencies. Unlike the current work, the analyses in these two papers do not include defensive investments. Nganje et al. (2008) also study the vulnerabilities in an interconnected system, focusing on a milk supply chain, where contamination can spread between nodes, and each node faces a binary decision of whether or not to invest in security measures.

In addition, the review article by Hausken and Levitin (2012) categorizes

articles by the defense measures used and by the attack tactics and circumstances. This research focuses on "protection" as the defense measure. The attack tactics and circumstances used are "attack against single element" and "random attack" (when the attacker is not strategic).

Unlike the current paper, the literature has typically looked at the concepts of coordination and decision making separately in different articles; most articles mentioned in this section examine a coordinated set of nodes or an uncoordinated set of nodes, but not both. One exception is Hausken (2002), which examines a combination of individual and collective interests in system reliability and discusses the possibility of coordination of investments under several different structures. The setting differs somewhat from the current paper; the defenders are protecting a public good, and there is no attacker choosing a single node to target.

Kunreuther and Heal (2003) develop several models for interdependent security investments where the choices are discrete, i.e. each agent chooses whether or not to make a given investment. This paper serves as a foundation for much of the other work referenced in this section. Zhuang et al. (2007) present equilibrium strategies for independent actors or groups where threats occur over time and investment decisions are binary. They do not model coordination explicitly, but allow for the use of subsidies. Hausken (2006a) examines interdependence and income effects in cybersecurity investments modeled as a competitive (uncoordinated) game between firms and an attacker, and includes a strategic attacker via a substitution effect of investment. Bakshi and Kleindorfer (2009) explore the challenge of independent investments in risk mitigation by two supply chain nodes using both non-cooperative and cooperative (bargaining) models. They observe underinvestment in the non-cooperative models, and obtain superior results with a cooperative approach; their work serves as a primary motivation for the present paper. Bandyopadhyay et al. (2010) consider network vulnerability in an analysis of cybersecurity investments by two supply chain partners. Based on their model, they also find that the firms tend to underinvest in cybersecurity. We expand upon this work by allowing for more than two firms and analyzing the problem using an attacker-defender model. Lee et al. (2011) explore collaboration between two supply chain partners in technology investment decisions, where security is one component of overall cost. They find that both underinvestment and overinvestment are possible when decisions are uncoordinated, depending on the relative importance of security and efficiency of the technology to the firms. Liu et al. (2011) consider a setting in which two firms can share security information with one another, and find that both underinvestment and overinvestment in security are possible in the absence of coordination, depending on the nature of the information.

Manshaei et al. (2013) provide a survey of work done in game theory and cybersecurity; one section provides several examples of papers that examine networks of interconnected nodes, which tend to involve utility functions that depend on the overall vulnerability of the network and/or measures of influence between nodes. Nagurney and Shukla (2017) use a measure of network vulnerability to capture dependence between nodes and interaction of investments, and

include a utility function over wealth levels.

It is important to note that network vulnerability is not the only way in which nodes can be interdependent in an attacker-defender model of supply chain cybersecurity. For example, if node 2 has node 1's customer information, then an attack on node 2 damages node 1 even if node 1's networks and physical assets are unaffected. Our model incorporates any indirect damage from an attack, irrespective of the underlying nature of the relationship between the nodes.

Attacker-defender models have been applied widely in the context of terrorism, where the defender is a nation or other decision making entity determining an optimal investment of resources. These models vary along many different dimensions, but tend to assume that the attacker is strategic. Baron et al. (2018) analyze both a single-period and a repeated game in which a strategic attacker has either one or multiple attack types, where the defender's choice is whether or not to respond. Zhuang and Bier (2007) present a model in which the attacker's effort level is endogenous, and may depend on defensive investments. Xu et al. (2016) consider a supply disruption to the defender and analyze possible supply chain risk management approaches; they find that while such strategies do generally benefit the defender, they can also increase the attacker's effort level. Hausken (2008) explores a model in which the terrorist is the defender, and a strategic attacker determines when to attack terrorist assets, and Hausken and Zhuang (2011) apply a model in which governments and terrorists are both attackers and defenders with endogenous levels of effort for both attacking and defending. Hausken (2017) models two interdependent targets, where each has a conditional probability of failure when the other fails. A single defender and a single attacker allocate effort to both targets, each of which is modeled as a contest. Hausken (2019) expands this analysis to more than two targets. Bier et al. (2007) explicitly consider the negative externality of pushing attack probability onto other defenders, and find that a centralized decision maker obtains better results than decentralized decision makers.

There is limited work, however, on indirect impacts of attacks across multiple defenders. One exception is Carceles-Poveda and Tauman (2011) who consider multiple governments collaborating to fight terrorists, though their model differs substantially from ours (as preemptive attacks are less accepted in the realm of private sector cybersecurity). Another is Kolfal et al. (2013), who analyze a model in which a firm's consumers react to IT security incidents occurring not only at that firm, but also at its competitors. They find that the nature of the elasticity of demand between firms when an incident occurs (complements, substitutes, or neither) can lead to different IT security spending equilibria.

3 Model and Analysis

Our model uses the following parameters:

n: the number of nodes in the supply chain

x_i: the amount of cybersecurity investment made by node i

$p(x_i)$: the attacker's probability of success if he attacks a node with cybersecurity investment x_i

d_{ij}: the damage to node j of a successful attack on node i, expressed in monetary units

The d_{ij} terms are what capture the supply chain relationships that are relevant to cybersecurity. For example, if nodes i and j share network infrastructure, we would expect d_{ij} and d_{ji} to be large. If node i shares its customer data with node j, we would expect d_{ji} to be large. High values of d_{ij} (where $i \neq j$) are analogous to a greater degree of interdependence in several of the papers discussed in Section 2.

We assume that $n \geq 2, x_i \geq 0$ for $i = 1, \ldots, n$, and $d_{ij} \geq 0$ for all pairs (i, j). In addition, we assume p is continuous, decreasing, and convex. These are common assumptions on attack probability of success functions; see, e.g., Bakshi and Kleindorfer (2009), Shetty et al. (2010), and Shan and Zhuang (2013a). However, there are many possible classes of p; see Hausken (2006b) for several examples (e.g., concave functions and logistic functions that are concave below a threshold and convex above it). The rationale behind the convexity assumption on p is that cybersecurity spending is prioritized based on cost-effectiveness, i.e. the most efficient investments are made first. Note that $p(x_i)$ does not depend on i; similar cybersecurity investment options are available to all nodes at comparable costs, and no node is implicitly easier or harder to attack than another.

In the initial model with a non-strategic attacker, we assume that the node to be attacked is chosen randomly and uniformally; that is, each node's probability of being attacked is $1/n$. Our model analyzes only a single attack; it is meant to be representative of attacks from a population of, initially, non-strategic attackers. The damage parameter is intentionally generic; it may include any one or more of the forms of cyberattack damage discussed in the Introduction.

3.1 Coordinated Supply Chain

In the ideal decision setting, all investment amounts are chosen by an optimal central planner. This is unlikely to be realistic for most supply chains, but is useful as a benchmark for optimal supply chain investment levels and results, and as a target for coordination mechanisms. With an optimal central planner, it is irrelevant which nodes bear which costs and damages; only the aggregate totals are relevant. Thus, it will be convenient to define an additional parameter:

$$d_i = \sum_{j=1}^{n} d_{ij}, \tag{1}$$

capturing the total damage resulting from a successful attack on node i.

The central planner's optimization problem is to minimize overall expected total cost, which can be expressed as:

$$\min_{x_1,\ldots,x_n} \sum_{i=1}^{n} x_i + \sum_{i=1}^{n} \frac{p(x_i)d_i}{n} \tag{2}$$

The first term of Equation (2) is simply the sum of the cybersecurity investment costs. The second term is the sum of the expected damages due to an attack for each node. Since Equation (2) is additively separable by node, it can be split into n individual optimization problems of the form:

$$\min_{x_i} x_i + \frac{p(x_i)d_i}{n}. \tag{3}$$

The first order condition of this optimization problem is:

$$1 + \frac{p'(x_i)d_i}{n} = 0 \tag{4}$$

and the second order condition (for convexity) is:

$$\frac{p''(x_i)d_i}{n} > 0 \tag{5}$$

Since p is assumed to be decreasing and convex, the second order condition is met, and the first order condition has a unique solution for x_i. Specifically, the optimal investment level x_i^* occurs where:

$$p'(x_i^*) = -\frac{n}{d_i}, \tag{6}$$

or:

$$x_i^* = \begin{cases} p'^{-1}\left(-\frac{n}{d_i}\right), & \text{if } d_i > p'(0)/n \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

Intuitively, the solution occurs when $p(x_i)$ has flattened enough such that the marginal investment cost is equal to the marginal reduction in expected damage. If d_i is small enough such that the marginal cost of investment exceeds the marginal reduction in expected damage at $x_i = 0$, then no positive value of x_i will satisfy the first-order condition, and $x_i^* = 0$. For most commonly used forms of p, the expected damage from an attack on node i at optimality is increasing in d_i; that is, even with optimal investments, the coordinated defender would still prefer that the attacker target the node i with the smallest value of d_i. We assume that this relationship between d_i and expected damage at optimality holds, as it does in the forms of p used by Shetty et al. (2010), Hausken and Zhuang (2013), Guan et al. (2017), and our example in Section 4. A sufficient condition for this relationship is: $p''(x_i) > p'(x_i)^2/p(x_i)$ for $i > 0$, which imposes a lower bound on the convexity of p.[1]

[1] One special case of interest is $p(x_i) = \alpha e^{-\beta x_i}$, where α, β are constants with $0 < \alpha \leq 1$ and $\beta > 0$; for this form of p, expected damage at optimality is fixed in our model and does not depend on d_i. This form has been used frequently in previous related work; see, e.g., Bier et al. (2008), Bakshi and Kleindorfer (2009), Shan and Zhuang (2013a), and Zhang et al. (2018).

The resulting expected total cost resulting from node i's investment and an attack on node i is:

$$\begin{cases} p'^{-1}\left(-\frac{n}{d_i}\right) + \frac{1}{n}p\left(p'^{-1}\left(-\frac{n}{d_i}\right)\right)d_i, & \text{if } d_i > p'(0)/n \\ \frac{p(0)d_i}{n}, & \text{otherwise} \end{cases} \tag{8}$$

Extending this result to the entire supply chain, the supply chain's expected cost C is given by:

$$C = \sum_{i=1}^{n}\left(p'^{-1}\left(-\frac{n}{d_i}\right) + \frac{1}{n}p\left(p'^{-1}\left(-\frac{n}{d_i}\right)\right)d_i\right), \tag{9}$$

if all d_i are large enough to justify positive investment levels. For any $x_i^* = 0$, the ith summand is $\frac{p(0)d_i}{n}$.

3.2 Uncoordinated Supply Chain

In this subsection, we consider the investment optimization of an independently-acting node i in the supply chain. In this case, the total damage d_i from an attack on node i does not affect the decision. Instead, the decision maker considers only d_{ii}: the damage to node i itself when it is attacked. The optimization is completely analogous to the individual node optimization in the coordinated supply chain case, with d_{ii} replacing d_i. While each node's decision might affect the damage experienced by the other nodes, it does not affect any of the other nodes' decision problems; thus, each one can be analyzed independently. Node i's optimization problem can be expressed as:

$$\min_{x_i} x_i + \frac{p(x_i)d_{ii}}{n}. \tag{10}$$

Theorem 1. *When the probabilities of each node being attacked are equal, and nodes are acting independently, each node will underinvest in cybersecurity relative to the supply chain optimal level (provided that level is positive).*

Proof. The unique optimal solution x_i^{**} for node i occurs where:

$$p'(x_i^{**}) = -\frac{n}{d_{ii}}, \tag{11}$$

or:

$$x_i^{**} = p'^{-1}\left(-\frac{n}{d_{ii}}\right). \tag{12}$$

Since $d_{ii} < d_i$, and p is convex, it is necessarily true that $x_i^{**} < x_i^*$. (Additionally, $d_{ii} < d_i$ implies that if $x_i^* = 0$, then $x_i^{**} = 0$ as well.) □

This result confirms that cybersecurity investment at each node will be lower if the nodes are acting independently, which is consistent with the results of Bakshi and Kleindorfer (2009) and Bandyopadhyay et al. (2010). In addition,

generally, the greater the disparity between d_{ii} and d_i (i.e. the indirect damages from an attack on node i), the greater the magnitude of underinvestment. Note also that it is possible for the ordering to change; i.e. that $x_i^* > x_j^*$ but $x_i^{**} < x_j^{**}$.

For the remainder of the paper, we assume for the sake of brevity that $x_i^*, x_i^{**} > 0$ for all i. This does not materially affect any results; an optimal investment level of zero indicates that the damage caused by an attack on that node is very small.

The resulting expected total cost (to the whole supply chain) resulting from node i's investment and an attack on node i is:

$$p'^{-1}\left(-\frac{n}{d_{ii}}\right) + \frac{1}{n}p\left(p'^{-1}\left(-\frac{n}{d_{ii}}\right)\right)d_i. \tag{13}$$

Since x_i^* is the investment level that minimizes expected cost to the supply chain, we know that (13) is greater than (8). The expected cost increase associated with node i's investment and attacks is:

$$\begin{aligned}\frac{1}{n}\left[p\left(p'^{-1}\left(-\frac{n}{d_{ii}}\right)\right) - p\left(p'^{-1}\left(-\frac{n}{d_i}\right)\right)\right]d_i \\ - \left[p'^{-1}\left(-\frac{n}{d_i}\right) - p'^{-1}\left(-\frac{n}{d_{ii}}\right)\right].\end{aligned} \tag{14}$$

The first term of (14) captures the increase in expected damage to the supply chain, and the second term captures the cost savings from reduced cybersecurity investment. Extending this result to the entire supply chain, the increase in total expected cost is:

$$\begin{aligned}\sum_{i=1}^{n}\Bigg(\frac{1}{n}\left[p\left(p'^{-1}\left(-\frac{n}{d_{ii}}\right)\right) - p\left(p'^{-1}\left(-\frac{n}{d_i}\right)\right)\right]d_i \\ - \left[p'^{-1}\left(-\frac{n}{d_i}\right) - p'^{-1}\left(-\frac{n}{d_{ii}}\right)\right]\Bigg).\end{aligned} \tag{15}$$

In the case of a non-strategic attacker, this sub-optimization can be avoided easily via coordination. The goal of coordination is to induce each node i to invest x_i^* in cybersecurity. The simplest approach is for all nodes to enter into a contract in which each node i must pay d_{ij} to every other node $j \neq i$ in the event of a successful attack on node i. When $n = 2$, this is roughly analogous to the contract proposed by Bandyopadhyay et al. (2010). Under such a contract, the uncoordinated optimization problem becomes equivalent to the coordinated optimization problem, because the actual cost to node i of being attacked will be d_i. Since this contract results in a net expected gain for the supply chain, it is possible for transfer payments to be established such that it is Pareto superior to the uncoordinated approach; indeed, this is the basis for bargaining approaches such as those used by Bakshi and Kleindorfer (2009).

Another possibility of note is that one firm might be willing to pay for a additional cybersecurity at a different firm. The effect of increasing x_i on firm

j's total expected cost is:

$$\frac{\partial}{\partial x_i}\left(\frac{p(x_i)d_{ij}}{n}\right) = \frac{p'(x_i)d_{ij}}{n} \tag{16}$$

At optimality, $x_i = x_i^{**} = p'^{-1}\left(-\frac{n}{d_{ii}}\right)$. Therefore, we can express the marginal effect on firm j's expected cost resulting from additional investment in firm i's cybersecurity as:

$$\frac{p'(x_i^{**})d_{ij}}{n}. \tag{17}$$

Combining this expression with Equation (11) allows us to state a helpful condition. Firm j is willing to invest unilaterally in additional cybersecurity for firm i if:

$$d_{ij} > d_{ii} \tag{18}$$

If this is the case for some pair of firms i and j, and firms are permitted to invest in other firms' cybersecurity, then the uncoordinated solution is in fact not even a Nash equilibrium. One could imagine this being the case if firm j is a large retailer and firm i is a small supplier. In such a situation, some coordination may occur even without a formal contractual mechanism, as occurred in the case of Aramco when they believed that current supplier cybersecurity was insufficient.

3.3 Strategic Attacker

In this subsection, we examine the impact of a strategic attacker who is more likely to attack nodes at which the expected total damage is greater. The standard game theory definition of a strategic attacker, in the context of this paper, is one who will choose the target that results in the greatest expected damage. We will discuss this type of strategic attacker briefly, but will focus more heavily on a *partially strategic* attacker who is more likely to attack nodes with greater expected damage, but whose choice of target cannot be predicted with certainty. We believe this formulation to be both more instructive and more realistic in this setting. There are certainly alternate ways to address the issue; for instance, the model used by Shan and Zhuang (2013b) includes a parameter representing the probability of the attacker being (fully) strategic.

The attack probability for node i is denoted by $a_i(p(x_1)d_1, \ldots, p(x_n)d_n)$. With a fully strategic attacker, if $p(x_i)d_i < \max_j p(x_j)d_j$, then $a_i = 0$. For a partially strategic attacker, we assume a_i is continuously differentiable, increasing in $p(x_i)d_i$, and decreasing in all $p(x_j)d_j, j \neq i$. For all $j \neq i$, we assume that a_i is concave in $p(x_j)d_j$ when $p(x_j)d_j < p(x_i)d_i$, and convex in $p(x_j)d_j$ when $p(x_j)d_j > p(x_i)d_i$. That is, the magnitude of impact on attack probability of a change to the expected damage at another node is smaller when the two expected damages are dissimilar. Figure 1 shows an illustration of this relationship. Finally, we assume that the attack probability functions are symmetric, in that results would not be affected by relabeling the nodes. None of these

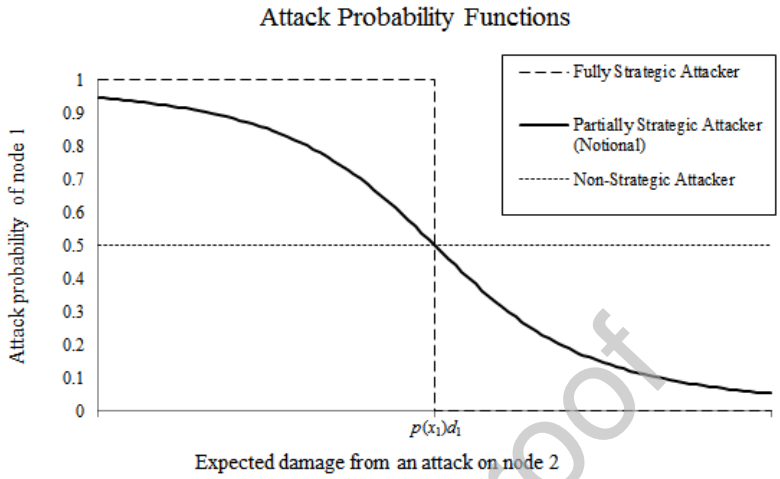

Figure 1: Attack probability functions for three types of attacker

assumptions is restrictive; we are simply asserting explicitly what it means for the attacker to be strategic.

For the sake of brevity, we will omit the vector of expected damages when expressing attack probabilities for the remainder of the paper, and will instead write simply a_i.

For a given vector of investment levels, node i's expected total cost is:

$$x_i + \sum_{j=1}^{n} a_j p(x_j) d_{ji}, \tag{19}$$

and the total expected cost to the supply chain is given by:

$$C = \sum_{j=1}^{n} x_j + \sum_{j=1}^{n} a_j p(x_j) d_j. \tag{20}$$

In addition, we can observe some straightforward implications of the assumptions on a_i:

- If $p(x_i)d_i = p(x_j)d_j$ for all i, j, then all of the attack probabilities must be $1/n$.
- If $p(x_i)d_i = p(x_j)d_j$ for any specific pair of nodes i, j, then $a_i = a_j$.
- If $p(x_i)d_i > p(x_j)d_j > p(x_k)d_k$, then $\frac{\partial a_j}{\partial x_i} > \frac{\partial a_k}{\partial x_i}$.

The third bullet point states that a change in the investment level of one node will have a greater impact on nodes whose expected damages are more similar.

3.3.1 Coordinated Supply Chain

In a coordinated supply chain, the central planner is optimizing the vector of investment levels x_i, similar to the case in which the attacker is non-strategic. However, the optimization problem can no longer be split into n independent optimizations, since the probability that node i will be attacked depends on the investment levels at the other nodes as well. Rather, the central planner must minimize (20) directly by the choice of $x_1, \ldots, x_n$.

With a fully strategic attacker, the set of decisions can be modeled as a Stackelberg (leader-follower) game, where the central planner is the leader, and the attacker is the follower. It is straightforward to observe that the investment levels should be chosen in such a way that the expected damages of attacking each node are equal (assuming, of course, that it is possible to do so). There would be no benefit to increasing investment at only a subset of nodes, since the attacker would simply not attack those nodes. Thus, the coordinated optimization problem with a fully strategic attacker amounts to choosing a level of expected damage d, and can be expressed as:

$$\min_d \; d + \sum_{i=1}^{n} p^{-1}\left(\frac{d}{d_i}\right); \tag{21}$$

that is, minimizing the expected damage plus the investments associated with it. Differentiating with respect to d yields the first-order condition:

$$1 + \sum_{i=1}^{n} \frac{1}{d_i} p^{-1'}\left(\frac{d}{d_i}\right) = 0. \tag{22}$$

Because p is decreasing and convex, p^{-1} is convex. Thus, total cost is convex in d, and there is a unique optimal solution.

The fully strategic case can be viewed as a limit; the more strategic the attacker is, the closer to equal the expected supply chain damages will be for an attack on each node.

The optimization problem for the partially strategic attacker is more involved, since it cannot be reduced to one decision variable. Differentiating the total cost function in (20) with respect to each investment level x_i yields a set of n first-order conditions:

$$a_i p'(x_i) d_i + 1 + \sum_{j=1}^{n} \frac{\partial a_j}{\partial x_i} p(x_j) d_j = 0. \tag{23}$$

These first-order conditions state that the marginal reduction in expected damage from an attack on any given node must be exactly offset by the marginal cost of the investment and the changes in expected damages due to the attacker's reaction.

Unfortunately, total cost is no longer convex in the investment levels; there can be many local minima, and there is no guarantee of a unique optimal solution. However, we can still identify informative properties that any optimal solution must satisfy, starting with the following two Lemmas:

Lemma 1. *The ordering of nodes from highest optimal investment level to lowest optimal investment level is the same with a strategic attacker as it is with a non-strategic attacker.*

The proof is given in the Appendix. The intuition behind this lemma is straightforward; regardless of whether the attacker is strategic or not, optimal investment level at a node is increasing in the amount of damage that would be caused by a successful attack on that node.

Lemma 2. *The ordering of nodes from highest optimal expected damage when attacked to lowest optimal expected damage when attacked is the same with a strategic attacker as it is with a non-strategic attacker (and is the same as the order from highest optimal investment level to lowest optimal investment level).*

The proof is given in the Appendix. Intuitively, if $d_i > d_j$, the marginal cost of reducing expecting damage at node i is higher than that at node j for any given amount of expected damage; thus, it will always be more cost-effective for expected damage from an attack on node j to be lower. As a reminder, $p(x_i)$ does not depend on i; if investment levels at two nodes are equal, then an attack on either is equally likely to be successful.

Given these results, we are able to state the following theorem:

Theorem 2. *The optimal coordinated cybersecurity investment levels have a wider range when facing a strategic attacker than when facing a non-strategic attacker.*

The proof is given in the Appendix. This is a somewhat counterintuitive result; one might imagine that a strategic attacker would force the coordinated defender to spread out resources more evenly between nodes (in an optimal solution) to avoid a weak link being exploited, but the opposite effect occurs. The reason for this is that the least appealing target for the attacker already has the lowest level of cybersecurity investment in the non-strategic case. At the margin, lowering this level further will transfer some attack probability from other nodes to the one at which an attack will do the least damage. By similar logic, investment will be increased at the node that is the most damaging target for the attacker in the non-strategic case. Figure 2 shows an illustrative notional result for an example with four nodes.

3.3.2 Uncoordinated Supply Chain

In the case of a fully strategic attacker when the supply chain is uncoordinated, there is generally no pure strategy equilibrium, as any given node i will prefer to set x_i such that $p(x_i)d_i$ is lower than, but infinitesimally close to, $\max_j p(x_j)d_j$. In practice, a fully strategic attacker is extremely unlikely, and the differences between a coordinated and uncoordinated supply chain given a strategic attacker can be understood without studying this boundary case. Therefore, we restrict our analysis in this section to a partially strategic attacker[2].

[2]There are also other modeling approaches that would lead to a pure strategy equilibrium. For instance, if some or all of the d_{ij} are uncertain when the investments are made, nodes

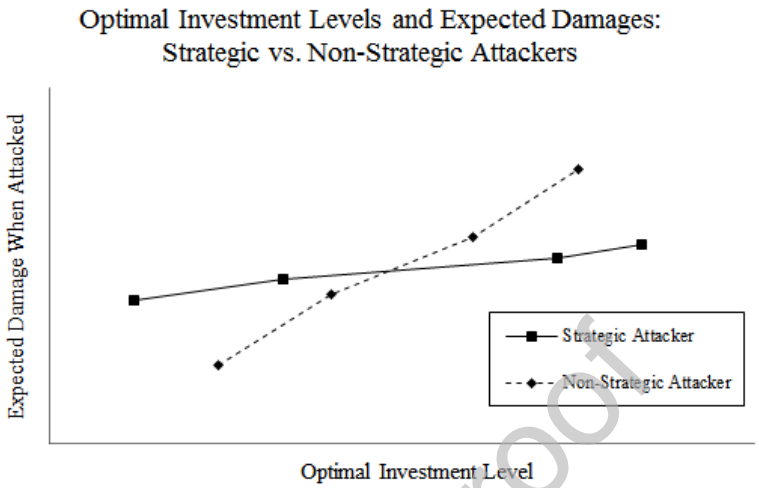

Figure 2: Supply chain optimal investment levels and expected damages for a notional four-node example

There are two factors that differentiate the uncoordinated problem from the coordinated problem when the attacker is strategic. The first is the same phenomenon observed with a non-strategic attacker: nodes only consider damages to themselves, not to other nodes. This leads to underinvestment. However, the second factor generally has the opposite effect: nodes gain additional benefit from their cybersecurity investments by reducing the probability that the attacker will target them. As we will observe, it is possible for either underinvestment or overinvestment to occur.

The individual node optimization problem with a strategic attacker can be expressed as:

$$\min_{x_i} x_i + \sum_{j=1}^{n} a_j p(x_j) d_{ji}. \tag{24}$$

Differentiating (24) with respect to x_i yields the first-order condition:

$$1 + a_i p'(x_i) d_{ii} + \sum_{j=1}^{n} \frac{\partial a_j}{\partial x_i} p(x_j) d_{ji} = 0. \tag{25}$$

We can think of (25) as capturing node i's best response function to a vector of levels of $x_j : j \neq i$. Thus, in equilibrium, (25) holds for all i.

The differences between (25) and the non-strategic case are that $a_i \neq 1/n$ in general, and that nodes can receive an additional benefit from investment

cannot know with certainty where a fully strategic attacker will attack; this is similar to the approach taken by Bier et al. (2007) regarding attacker preferences.

by transferring attack probability to other nodes. Unfortunately, as in the coordinated problem, a node's expected damage is not convex in its investment level, which makes it difficult to identify one best response, let alone an equilibrium set of investment levels. However, we provide two specific results that are helpful for understanding equilibrium behavior with a strategic attacker:

Theorem 3. *If $d_{j'i'} = 0$ for all $i', j' : i' \neq j'$, then for any optimal set of investment levels in the coordinated case, there exists a pure strategy equilibrium in which every node overinvests.*

That is, if there are no indirect damages, a strategic attacker leads the nodes to invest *more* than they would when coordinating. The proof is given in the Appendix. Intuitively, when nodes are acting independently, they ignore the cost to the supply chain of pushing attack probability onto other nodes.

Theorem 4. *If $d_{j'i'} = d_{j'}/n$ for all i', j', then for any optimal set of investment levels in the coordinated case, there exists a pure strategy equilibrium in which every node underinvests.*

This result occurs because when damages are split equally, the best response conditions are equivalent to the coordinated case first order conditions, except that all damages are divided by n. The proof is in the Appendix; it consists primarily of showing that for a set of optimal coordinated investment levels x_i^*, if damages are divided by n, a corresponding equilibrium exists with $x_i < x_i^*$ for all i.

Thus, when there are no indirect damages, a pure strategy equilibrium exists in which all nodes overinvest, and when indirect damages are equal in magnitude to direct damages, a pure strategy equilibrium exists in which all nodes underinvest. For intermediate or disparate levels of indirect damages, it is possible that some nodes will overinvest and others will underinvest.

It is also important to note that a simple coordination mechanism is no longer possible with a strategic attacker. There is no guarantee of a unique optimal solution to the coordinated problem. Therefore, no general transfer payment strategy can ensure that all nodes will choose investment levels that are part of the same optimal solution. It is paramount that the organizations involved work closely together to ensure investment levels that are not detrimental to the supply chain. While this is certainly a significant challenge, it can be addressed in many different ways through improved relationships with suppliers (Rinehart et al., 2004).

3.4 Supply Chain Considerations

To keep the analytical results of the paper as general as possible, we have not imposed any structure or restrictions on the set of d_{ij} to this point. However, their purpose is to capture the relationships between various nodes in the supply chain, and it is therefore worthwhile to understand the impacts of characteristics that are likely to arise in practice. In this section, we explore two such traits that are common in supply chains.

The first common trait is that indirect damages of an attack tend to be higher for nodes that are nearer to the attacked node in the supply chain than nodes that are farther removed. This would suggest that d_{ij} decreases in $|i-j|$, all else being equal. Consider a supply chain in which $d_{ij} = f(|i-j|) \geq 0$ for all i, j, where f is a monotonically decreasing function. That is, indirect damage is decreasing in (and determined entirely by) the distance between the two nodes, and is lower than direct damage, which is equal for every node.

It is straightforward to observe that d_i is highest at $i = \frac{n+1}{2}$ if n is odd, and at $i = \frac{n}{2}$ and $i = \frac{n}{2} + 1$ if n is even (i.e. at the middle node(s)), and that d_i decreases symmetrically toward the ends of the supply chain. With a non-strategic attacker, it follows immediately from Equation 12 that if nodes do not coordinate, all of their optimal investment levels will be equal, since d_{ii} is equal for all i. From Equation 7, however, the coordinated solution involves investments that are increasing in d_i, and thus are highest for the nodes in the middle of the supply chain and lowest at the ends. While all individual nodes underinvest in the absence of coordination, it is most pronounced in the middle of the supply chain.

If the attacker is strategic, the coordinated solution still involves investments that are highest in the middle of the supply chain and lowest at the ends (Lemma 1), but the range of these investments will be wider than in the case of a non-strategic attacker (Theorem 2). That is, the result of investments being concentrated toward the middle of the supply chain is even more pronounced.

When the attacker is strategic and nodes are uncoordinated, it is possible for multiple equilibria to exist, and both overinvestment and underinvestment can occur. The general forms of f and a_i preclude a more specific result. However, Theorems 3 and 4 are instructive here. If $f(|i-j|)$ decreases substantially in $|i-j|$, i.e. indirect damages are small relative to direct damages, we would expect to observe overinvestment. If f is relatively flat as $|i-j|$ increases, we would expect to observe underinvestment.

The second common trait that might arise in practice is that damages flow more easily or more substantially in one direction than the other. This could be due to certain types of information, such as orders or customer data, flowing in one direction, or to goods flowing in one direction. This asymmetry of indirect damages can be modeled in several ways. The approach chosen here is intended to isolate the effect. We impose the condition that $d_{ij} = d_h$ for $i < j$, and $d_{ij} = d_l$ for $i > j$, where $d_h > d_l \geq 0$. That is, every indirect damage in one direction (toward higher numbered nodes) is equal to a constant, and every indirect damage in the other direction is equal to a smaller constant. In addition, we assume d_{ii} does not depend on i; that is, direct damage is equal for every node. Note that these two conditions together imply that d_i is decreasing in i. It is straightforward to adapt this approach to damages that flow in the opposite direction.

As before, since d_{ii} is equal for all i, it is clear from Equation 12 that if the attacker is non-strategic and nodes do not coordinate, all of their optimal investment levels will be equal. Additionally, since d_i is decreasing in i, Equation 7 implies that optimal coordinated investment levels will be decreasing in i as

Table 1: Coordinated and uncoordinated solutions: non-strategic attacker

Coordinated Solution			Uncoordinated Solution			
Node	x_i^*	$p(x_i^*)d_i$	Node	x_i^*	$p(x_i^*)d_{ii}$	$p(x_i^*)d_i$
1	4.77	17.32	1	3.83	14.49	20.70
2	2.65	10.95	2	1.23	6.71	17.89
3	1.00	6.00	3	0.63	4.90	7.35

well; this is consistent with the results of Hausken (2019). Without coordination, all nodes underinvest; the magnitude of underinvestment is highest at the lowest-numbered nodes (where an attack causes the most indirect damage).

The effect of the attacker being strategic is the same for the coordinated solution as observed previously: the ordering of nodes by investment level is preserved, but the range of investment levels becomes wider. In this case, that means higher investments at the lowest-numbered nodes, and lower investments at the higher-numbered nodes. In the uncoordinated case, as previously, there may be multiple equilibria; specific results will depend on the relative levels of d_h, d_l, and d_{ii}, and the form of a_i, but Theorems 3 and 4 are again instructive. Because the disparity between d_i and d_{ii} is largest for the lowest-numbered nodes, we would tend to observe more underinvestment at those nodes, and more overinvestment at the highest-numbered nodes.

4 Example

In this section, we consider a three-node numerical example to illustrate some of the important concepts and relationships explored in the previous sections. We let $p(x_i) = 1/(x_i + 1)$. The following matrix specifies all nine values of d_{ij} (row i and column j):

$$\begin{bmatrix} 70 & 20 & 10 \\ 20 & 15 & 5 \\ 8 & 1 & 3 \end{bmatrix}$$

This matrix implies that $d_1 = 100$, $d_2 = 40$, and $d_3 = 12$, and that node 1 is the most severely impacted, regardless of which node is actually attacked. This is meant to reflect the case in which node 1 is a large retailer, and nodes 2 and 3 are smaller nodes earlier in the supply chain.

Table 1 shows the coordinated and uncoordinated solutions when the attacker is non-strategic. The coordinated solution was obtained using Equation (7), and the uncoordinated solution was obtained using Equation (12). Note that the uncoordinated investments are far lower than the coordinated investments. Total cybersecurity spending is 2.73 lower in the uncoordinated case, but expected damage to the supply chain is 3.89 higher.

To examine the impact of a strategic attacker, consider attack probabilities a_i defined as follows. Let ϵ_i denote a uniform random variable with support

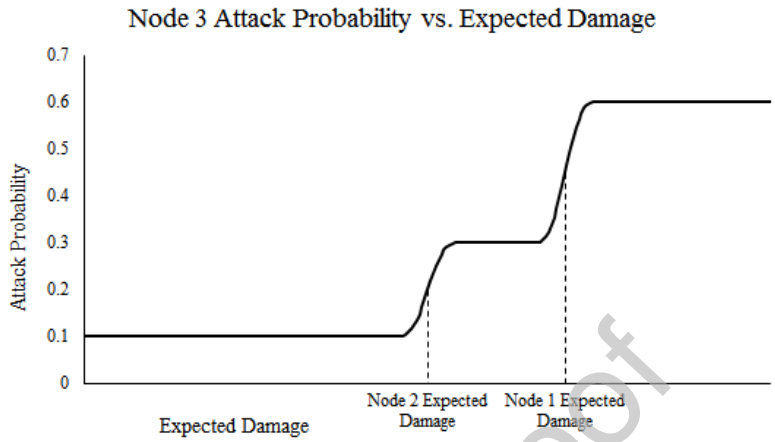

Figure 3: Change in attack probability of node 3 as its expected damage increases

$[-0.1, 0.1]$, and let r_i denote node i's rank (between 1 and n) when nodes are ordered by $p(x_i)d_i + \epsilon_i$. That is, the nodes are ranked by expected damage, subject to an error term; $r_i = 1$ represents the node with the highest value of $p(x_i)d_i + \epsilon_i$. Each rank has an attack probability associated with it, and the error terms are not realized until after the investment levels have been chosen. In this example, the highest ranked node will be attacked with probability 0.6, the second-highest with probability 0.3, and the lowest with probability 0.1. An illustration of the attack probability function for node 3's expected damage (given fixed expected damages for nodes 1 and 2) is shown in Figure 3.

Table 2 shows the coordinated solution and an uncoordinated pure strategy equilibrium for these three nodes given this particular strategic attacker. The coordinated solution was obtained using Equation (23); in this case, there is a unique optimal solution. The uncoordinated equilibrium was determined using Equation (25). Note that there is no guarantee that this is the only pure strategy equilibrium. In addition, it assumes common knowledge for all decision makers, which is a particularly strong assumption when a strategic attacker's behavior is included in the model. Total investment level is 0.71 lower in the uncoordinated case, while expected damage to the supply chain is 0.77 higher.

Two critical observations can be made by comparing Tables 1 and 2. First, if we compare the two coordinated solutions, we can see that the optimal investment levels are much more disparate with a strategic attacker, and the resulting expected damages are much more similar to one another. Intuitively, this occurs because the strategic attacker is more likely to target nodes at which expected damage is greater; hence, there is greater benefit to increasing investment at those nodes (and vice versa for nodes at which expected damage is low).

Second, we can observe that the uncoordinated investment levels are much

Table 2: Coordinated and uncoordinated solutions: strategic attacker

Coordinated Solution			Uncoordinated Solution			
Node	x_i^*	$p(x_i^*)d_i$	Node	x_i^*	$p(x_i^*)d_{ii}$	$p(x_i^*)d_i$
1	6.75	12.91	1	6.54	9.28	13.26
2	2.46	11.55	2	2.06	4.90	13.07
3	0.10	10.95	3	0.00	8.00	12.00

closer to the coordinated investment levels when the attacker is strategic. Each node still underinvests, but to a lesser degree. This is due to the overinvestment effect of a strategic attacker; a node acting independently can gain additional benefit from transferring attack probability to other nodes.

5 Conclusion

Coordination of cybersecurity investments across firms in a supply chain is critical, as attacks on any one nodes will impact the rest of the supply chain as well. We have developed a model to analyze both optimal cybersecurity investment levels for coordinated supply chains and the cybersecurity investment levels resulting from nodes acting independently. We obtained results both for a non-strategic and strategic attacker, and obtained several important insights. In all cases, the analysis was carried out for arbitrarily long supply chains; the model is not limited to two nodes. These analytical results were then demonstrated using a numerical example.

First, consistent with prior work, we found that when the attacker is non-strategic, nodes acting independently will underinvest in cybersecurity relative to the supply chain optimal levels. Intuitively, this is due to independent nodes ignoring the indirect damages to the rest of the supply chain when determining their investment levels. Simple coordination mechanisms can induce supply chain optimal cybersecurity investment decisions for each node.

Second, we observed that even in the absence of a formal coordination mechanism, larger nodes might choose to invest in or subsidize smaller nodes' cybersecurity. This occurs if indirect damages to the larger node are greater than direct damages to the smaller node when the smaller node is attacked.

Third, we found that when the attacker is strategic, the supply chain optimal investment levels are more disparate, while the resulting expected damages from attacks on each node are more similar to one another. That is, with a strategic attacker, it becomes desirable to increase spending at the larger nodes and decrease spending at the smaller nodes, making the attacker closer to indifferent regarding the choice of target.

Finally, the impacts of a lack of coordination and a strategic attacker somewhat counterbalance one another. A lack of coordination leads to underinvestment in cybersecurity, but independently-acting nodes will invest more when the attacker is strategic than when the attacker is non-strategic (due to the added

benefit of pushing attack probability onto other nodes). Therefore, it is possible to observe either underinvestment or overinvestment in an uncoordinated supply chain when the attacker is strategic. Overinvestment is more likely to occur when indirect damages of attacks are very low relative to direct damages.

The interactions between supply chain coordination and attacker behavior are complex, and certainly merit additional study. For example, future work might explore a network of nodes rather than a single supply chain, where a cyberattack on a single node impacts nodes in multiple supply chains, and a strategic attacker has a broader set of choices.

It might also be fruitful to analyze a portfolio-based model in which cybersecurity investments are discrete rather than continuous. However, this would require qualitatively different approaches and may yield very different insights, such as potential benefit of keeping investments secret (Dighe et al., 2009), as the analysis in the current paper relies on convexity properties of the continuous functions involved.

References

AIG (2013). Cyber and data security risks and the real estate industry. Last modified December 17, 2018. `https://www.aig.com/content/dam/aig/america-canada/us/documents/business/industry/aig-cyber-edge-whitepaper-final-brochure.pdf`.

Allianz Global Corporate and Specialty (2017). Allianz risk barometer: Top business risks 2017. Last modified December 17, 2018. `https://www.allianz.at/v_1484002800000/ueber-allianz/media-newsroom/news/aktuelle-news/pa-download/20170111allianz-risk-barometer-2017-report.pdf`.

Bakshi, N. and Kleindorfer, P. (2009). Co-opetition and investment for supply-chain resilience. *Production and Operations Management*, 18(6):583–603.

Bandyopadhyay, T., Jacob, V., and Raghunathan, S. (2010). Information security in networked supply chains: impact of network vulnerability and supply chain integration on incentives to invest. *Information Technology and Management*, 11(1):7–23.

Baron, O., Berman, O., and Gavious, A. (2018). A game between a terrorist and a passive defender. *Production and Operations Management*, 27(3):433–457.

Bier, V., Haphuriwat, N., Menoyo, J., Zimmerman, R., and Culpen, A. (2008). Optimal resource allocation for defense of targets based on differing measures of attractiveness. *Risk Analysis*, 28(3):763–770.

Bier, V., Oliveros, S., and Samuelson, L. (2007). Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*, 9(4):563–587.

Boyson, S. (2014). Cyber supply chain risk management: Revolutionizing the strategic control of critical it systems. *Technovation*, 34(7):342–353.

Carceles-Poveda, E. and Tauman, Y. (2011). A strategic analysis of the war against transnational terrorism. *Games and Economic Behavior*, 71(1):49–65.

Chopra, S. S. and Khanna, V. (2015). Interconnectedness and interdependencies of critical infrastructures in the us economy: Implications for resilience. *Physica A: Statistical Mechanics and its Applications*, 436:865–877.

Davis, A. (2015). Building cyber-resilience into supply chains. *Technology Innovation Management Review*, 5(4).

Dighe, N. S., Zhuang, J., and Bier, V. (2009). Secrecy in defensive allocations as a strategy for achieving more cost-effective attacker deterrence. *International Journal of Performability Engineering*, 5(1):31–43.

Gaudenzi, B. and Siciliano, G. (2017). Just do it: Managing it and cyber risks to protect the value creation. *Journal of Promotion Management*, 23(3):372–385.

Guan, P., He, M., Zhuang, J., and Hora, S. (2017). Modeling a multi-target attacker-defender game with budget constraints. *Decision Analysis*, 14(2):87–107.

Hausken, K. (2002). Probabilistic risk analysis and game theory. *Risk Analysis*, 22(1):17–27.

Hausken, K. (2006a). Income, interdepenendence, and substitution effects affecting incentives for security investment. *Journal of Accounting and Public Policy*, 25(6):629–665.

Hausken, K. (2006b). Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability. *Information Systems Frontiers*, 8(5):338–349.

Hausken, K. (2008). Whether to attack a terrorist's resources stock today or tomorrow. *Games and Economic Behavior*, 64:548–564.

Hausken, K. (2017). Defense and attack for interdependent systems. *European Journal of Operational Research*, 256(2):582–591.

Hausken, K. (2019). Defence and attack of complex interdependent systems. *Journal of the Operational Research Society*, 70(3):364–376.

Hausken, K. and Levitin, G. (2012). Review of systems defense and attack models. *Journal of Performability Engineering*, 8(4):355–366.

Hausken, K. and Zhuang, J. (2011). Governments' and terrorists' defense and attack in a t-period game. *Decision Analysis*, 8(1):46–70.

Hausken, K. and Zhuang, J. (2013). The impact of disaster on the interaction between company and government. *European Journal of Operational Research*, 225(2):363–376.

Khan, O. and Estay, D. A. S. (2015). Supply chain cyber-resilience: Creating an agenda for future research. *Technology Innovation Management Review*, April:6–12.

Kolfal, B., Patterson, R., and Yeo, L. (2013). Market impact on it security spending. *Decision Sciences*, 44(3):517–556.

Kumar, R. and Park, S. (2019). A portfolio approach to supply chain risk management. *Decision Sciences*. Forthcoming.

Kunreuther, H. and Heal, G. (2003). Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249.

Lee, J., Palekar, U. S., and Qualls, W. (2011). Supply chain effiiciency and security: Coordination for collaborative investment in technology. *European Journal of Operational Research*, 210(3):568–578.

Liu, D., Ji, Y., and Mookerjee, V. (2011). Knowledge sharing and investment decisions in information security. *Decision Support Systems*, 52(1):95–107.

Mahwah, N. J. (2017). Radware research finds data loss is top cyber-attack concern. Last modified December 17, 2018. `https://www.radware.com/newsevents/pressreleases/2017/ert2016-2017/`.

Manshaei, M. H., Zhu, Q., Alpcan, T., Basar, T., and Hubaux, J. P. (2013). Game theory meets networks security and privacy. *ACM Computing Surveys*, 45(3):588–600.

McAfee Center for Strategic and International Studies (2014). Net losses: Estimating the global cost of cyber-crime. Last modified December 17, 2018. `https://www.mcafee.com/enterprise/en-us/solutions/lp/economics-cybercrime.html`.

Modi, S. B., Wiles, M. A., and Mishra, S. (2015). Shareholder value implications of service failures in triads: The case of customer information security breaches. *Journal of Operations Management*, 35:21–39.

Nagurney, A. and Shukla, S. (2017). Multifirm models of cybersecurity investment competition vs. cooperation and network vulnerability. *European Journal of Operational Research*, 260(2):588–600.

Nganje, W., Bier, V., Han, H., and Zack, L. (2008). Models of interdependent security along the milk supply chain. *American Journal of Agricultural Economics*, 90(5):1265–1271.

Ramakrishnan, S. and Bose, N. (2017). Target in $18.5 million multi-state settlement over data breach. Last modified January 10, 2019. https://www.reuters.com/article/us-target-cyber-settlement/target-in-18-5-million-multi-state-settlement-over-data-breach-idUSKBN18J2GH.

Rao, S. and Goldsby, T. J. (2009). Supply chain risks: A review and typology. *The International Journal of Logistics Management*, 20(1):97–123.

Rinehart, L. M., Myers, M. B., and Eckert, J. A. (2004). Supplier relationships: the impact on security. *Supply Chain Management Review*, 8(6):52–59.

Shan, X. and Zhuang, J. (2013a). Cost of equity in homeland security resource allocation in the face of a strategic attacker. *Risk Analysis*, 33(6):1083–1099.

Shan, X. and Zhuang, J. (2013b). Hybrid defensive resource allocations in the face of partially strategic attackers in a sequential defender–attacker game. *European Journal of Operational Research*, 228(1):262–272.

Shetty, N., Schwartz, G., Felegyhazi, M., and Walrand, J. (2010). Competitive cyber-insurance and internet security. In Moore, T., Pym, D. J., and Ioannidis, C., editors, *Economics of Information Security and Privacy*, pages 229–247. Springer, New York.

Smith, M. (2016). Huge rise in hack attacks as cyber-criminals target small businesses. *The Guardian*. Last modified January 3, 2019. https://www.theguardian.com/small-business-network/2016/feb/08/huge-rise-hack-attacks-cyber-criminals-target-small-businesses.

Williams, C. (2014). Security in the cyber supply chain: Is it achievable in a complex, interconnected world? *Technovation*, 34(7):382–384.

Woods, B. and Bochman, A. (2018). Supply chain in the software era. *World Economic Forum*. Last modified January 2, 2019. https://www.weforum.org/agenda/2018/06/managing-risk-in-the-energy-sector-s-cyber-supply-chain/.

Wu, B., Tang, A., and Wu, J. (2016). Modeling cascading failures in interdependent infrastructures under terrorist attacks. *Reliability Engineering and System Safety*, 147:1–8.

Xu, J., Zhuang, J., and Liu, Z. (2016). Modeling and mitigating the effects of supply chain disruption in a defender–attacker game. *Annals of Operations Research*, 236(1):255–270.

Yang, C.-C. and Wei, H.-H. (2013). The effect of supply chain security management on security performance in container shipping operations. *Supply Chain Management: An international Journal*, 18(1):74–85.

Zhang, J., Zhuang, J., and Jose, V. R. R. (2018). The role of risk preferences in a multi-target defender-attacker resource allocation game. *Reliability Engineering and System Safety*, 169:95–104.

Zhuang, J. and Bier, V. (2007). Balancing terrorism and natural disasters - defensive strategy with endogenous attacker effort. *Operations Research*, 55(5):976–991.

Zhuang, J., Bier, V., and Gupta, A. (2007). Subsidies in interdependent security with heterogeneous discount rates. *The Engineering Economist*, 52(1):1–19.

6 Appendix

Proof of Lemma 1. Consider two arbitrary distinct nodes i, j for which $d_i > d_j$. In the coordinated non-strategic case with optimal investment levels x_i^*, x_j^*, we know that $x_i^* > x_j^*$ and $p(x_i^*)d_i > p(x_j^*)d_j$. To show that the ordering of nodes by optimal investment level is identical for the strategic and non-strategic cases, it will suffice to show that a set of investment levels with $x_j > x_i$ cannot be optimal in the strategic case.

Total supply chain cost in the strategic case, as given by Equation (20), is:

$$C = \sum_{j=1}^{n} x_j + \sum_{j=1}^{n} a_j p(x_j) d_j.$$

Denoting the expected damage from an attack on node i as z_i, and the cost of obtaining it as $c_i(z_i)$ we can rewrite the total supply chain cost as:

$$C = \sum_{j=1}^{n} c_i(z_i) + \sum_{j=1}^{n} a_j z_j.$$

First, we consider the case where $x_j > x_i$ (which implies $z_i > z_j$). In this case, it is possible to decrease total supply chain cost simply by swapping z_i and z_j; that is, by investing x_i' and x_j' such that $p(x_i')d_i = z_j$ and $p(x_j')d_j = z_i$. (If x_j is high enough such that z_i cannot be reduced to z_j, and/or x_i is low enough such that z_j cannot be raised to z_i, then it is trivial to observe that the set of investment levels is not optimal.)

By symmetry of a, swapping z_i and z_j will swap the attack probabilities of nodes i and j, and will not change the attack probability at any other node. (And, of course, it will not change any of the other investment levels.) All that is required is to show that:

$$c_i(z_j) + c_j(z_i) < c_i(z_i) + c_j(z_j)$$

Rearranging this expression, we obtain:

$$c_i(z_j) - c_j(z_j) < c_i(z_i) - c_j(z_i)$$

Because c_i is simply $p^{-1}(z/d_i)$, we can observe that $c_i - c_j$ (i.e. the cost difference between the two nodes for achieving an expected damage level z)

is increasing in z, and therefore the inequality holds; total cost is lower when the expected damages are swapped.

Thus, we have found a set of investment levels with a lower total supply chain cost; a set of investment levels with $x_j > x_i$ cannot be optimal, which establishes the Lemma. □

Proof of Lemma 2. Total supply chain cost in the strategic case, as given by Equation (20), is:

$$C = \sum_{j=1}^{n} x_j + \sum_{j=1}^{n} a_j p(x_j) d_j.$$

Denote the expected damage from an attack on node i as z_i, and consider two nodes i and j such that in an optimal set of investment levels with a strategic attacker, z_i and z_j are adjacent when expected damages are ordered from highest to lowest. Arbitrarily, let $d_i > d_j$, which implies $z_i > z_j$ in the non-strategic optimum. From Lemma 1, $x_i > x_j$ in the strategic optimum.

The derivative of total supply chain cost with respect to x_i is:

$$a_i p'(x_i) d_i + 1 + \sum_{k=1}^{n} \frac{\partial a_k}{\partial x_i} p(x_k) d_k.$$

We can rewrite it as:

$$1 + \frac{\partial z_i}{\partial x_i} \left(a_i + \sum_{k=1}^{n} \frac{\partial a_k}{\partial z_i} p(x_k) d_k \right),$$

and similarly for x_j. Let $z_i = z_j$. (This implies $x_i > x_j$, and $a_i = a_j$.) The term in parentheses is then equal for nodes i and j, by symmetry of a. Based on the non-strategic optimal solution, it is clear that the partial derivative of z with respect to x will be higher (less negative) for node i than node j. Thus, the marginal change in supply chain cost will be farther away from 1 for node j than node i. That implies that if it is beneficial at the margin to increase the investment level at node i (to obtain a lower expected damage than at node j), it is even more beneficial to increase the investment level at node j.

This disparity becomes even larger as x_i increases further. Since a_i and p are both decreasing in magnitude as x_i increases, the marginal reduction in expected damage from an attack on node i is decreasing, and by convexity/concavity properties of a, more attack probability is being shifted to the more damaging nodes, and less to the less damaging nodes. Thus, it is impossible for both first order conditions to be satisfied when $z_i \leq z_j$, and we can conclude that $z_i > z_j$ in any optimal solution. Since i and j were arbitrarily chosen nodes adjacent in ordering by expected damage, repeated application of this result establishes the Lemma. □

Proof of Theorem 2. Recall that when the attacker is non-strategic, the attack probabilities are $a_i = 1/n$ for all i. Consider the node l with the lowest value of d_i, and thus lowest optimal investment level x_i^* in the non-strategic case. From

Lemmas 1 and 2, node l must also have the lowest optimal investment level and expected damage in the strategic case.

To establish the Theorem, we will show that, in the strategic case, total supply chain cost is increasing in x_l on the interval $[x_l^*, x_k]$, where x_k is the next lowest investment level. (An analogous argument can be made for the node h with the highest value of d_i.)

The derivative of total supply chain cost with respect to x_l is:

$$a_l p'(x_l) d_l + 1 + \sum_{j=1}^{n} \frac{\partial a_j}{\partial x_l} p(x_j) d_j.$$

Since a_l is the lowest attack probability, it is necessarily less than $1/n$, and thus the first term of the derivative with respect to supply chain cost is greater than -1 when $x_l = x_l^*$. This term is increasing in x_l, and therefore is greater than -1 on $[x_l^*, x_k]$. The third term (the marginal impact of changes in attack probabilities) is positive as long as $x_l \leq x_k$, since the smallest $p(x_j)d_j$ occurs when $j = l$, the partial derivatives of the a_j must sum to zero, and the partial derivative of a_l is the only one that is negative. Thus, the marginal impact on supply chain cost of increasing x_l is positive on $[x_l^*, x_k]$, and therefore an optimal set of investment levels with a strategic attacker must have $x_l < x_l^*$. □

Proof of Theorem 3. The best response condition for node i is:

$$1 + a_i p'(x_i) d_{ii} + \sum_{j=1}^{n} \frac{\partial a_j}{\partial x_i} p(x_j) d_{ji} = 0.$$

Since $d_{ji} = 0$ for all $j \neq i$, the condition reduces to:

$$1 + a_i p'(x_i) d_{ii} + \frac{\partial a_i}{\partial x_i} p(x_i) d_{ii} = 0.$$

We can compare this to the first order condition in the coordinated optimization (also with $d_{ji} = 0$ for all $j \neq i$, meaning that $d_i = d_{ii}$ for all i):

$$1 + a_i p'(x_i) d_{ii} + \sum_{j=1}^{n} \frac{\partial a_j}{\partial x_i} p(x_j) d_{jj} = 0.$$

The left hand side of this condition is strictly higher, since it includes the increase in expected damage from shifting attack probability to other nodes (but is otherwise equal). Therefore, the left hand side of the best response condition will be negative for any $x_1^*, \ldots, x_n^*$ satisfying the coordinated first order conditions; that is, a marginal increase in any node's investment level will decrease its expected cost.

By continuity of p and the attack probability functions, each node's expected cost function is also continuous. It is also clearly increasing for sufficiently large investment levels, meaning there must exist a set of best responses with $x_i > x_i^*$ for all i. Any such local optimum is also a set of best responses in the uncoordinated problem, which establishes the Theorem. □

Proof of Theorem 4. The first derivative of node i's cost function is:

$$1 + a_i p'(x_i) d_{ii} + \sum_{j=1}^{n} \frac{\partial a_j}{\partial x_i} p(x_j) d_{ji}.$$

When $d_{ji} = d_{jj}$ for all i, j, this derivative can be simplified to:

$$1 + a_i p'(x_i) \frac{d_i}{n} + \sum_{j=1}^{n} \frac{\partial a_j}{\partial x_i} p(x_j) \frac{d_j}{n},$$

Note that this is identical to the first derivative of total supply chain cost in node i's investment level in the coordinated case, but with all damages divided by n. Thus, the first order conditions from that coordinated optimization problem are equivalent to the best response conditions here; any locally optimal coordinated solution will be a pure strategy equilibrium for this problem. Consider a set of investment levels $x_1^*, \ldots, x_n^*$ that is supply chain optimal for damages $d_1, \ldots, d_n$. To establish the Lemma, we will show that for damages $d_1/n, \ldots, d_n/n$, there exists an optimal set of investment levels $x_1, \ldots, x_n$ such that $x_i < x_i^*$ for all i.

Expanding the derivative of the attack probabilities in the previous expression yields:

$$1 + \frac{1}{n} p'(x_i) d_i \left(a_i + \sum_{j=1}^{n} \frac{\partial a_j}{\partial z_i} p(x_j) d_j \right),$$

When $x_i = x_i^*$ for all i, this derivative is equal to $1 - 1/n$. That is, at $x_1^*, \ldots, x_n^*$, total cost is increasing in x_i for all i; a marginal decrease in any investment level will lead to a lower total supply chain cost. By continuity of p and the attack probability functions, the supply chain cost function is also continuous. It is also clearly bounded for finite investment levels, meaning there must exist a local optimum with $x_i < x_i^*$ for all i. Any such local optimum is also a set of best responses in the uncoordinated problem, which establishes the Theorem. □