Received: 16 January 2022 | Accepted: 21 June 2023

DOI: 10.1111/poms.14047

POMS

ORIGINAL ARTICLE

Cooperative security against interdependent risks

Sanjith Gopalakrishnan[1] | **Sriram Sankaranarayanan**[2]

[1]Desautels Faculty of Management, McGill University, Montreal, Quebec, Canada

[2]Indian Institute of Management Ahmedabad, Ahmedabad, India

Correspondence
Sanjith Gopalakrishnan, Desautels Faculty of Management, McGill University, Montreal, Quebec H3A 1G5, Canada.
Email: sanjith.gopalakrishnan@mcgill.ca

Handling Editor: Jayashankar Swaminathan

Abstract

Firms in interorganizational networks are exposed to interdependent risks that are transferable across partner firms, such as contamination in food supply chains or data breaches in technology networks. They can be decomposed into intrinsic risks a firm faces from its own operations and extrinsic risks transferred from its partners. Firms have access to two security strategies: either they can independently eliminate both intrinsic and extrinsic risks by securing their links with partners or, alternatively, firms can cooperate with partners to eliminate sources of intrinsic risk in the network. We develop a graph-theoretic model of interdependent security and demonstrate that the network-optimal security strategy can be computed in polynomial time. Then, we use cooperative game-theoretic tools to examine, under different informational assumptions, whether firms can sustain the network-optimal security strategy via suitable cost-sharing mechanisms. We design a novel cost-sharing mechanism: a restricted variant of the well-known Shapley value, the *agreeable allocation*, that is easy to compute, bilaterally implementable, ensures stability, and is fair. However, the agreeable allocation need not always exist. Interestingly, we find that in networks with homogeneous cost parameters, the presence of locally dense clusters of connected firms precludes the existence of the agreeable allocation, while the absence of sufficiently dense clusters (formally, k-cores) guarantees its existence. Finally, using the SDC Platinum database, we consider all interfirm alliances formed in the food manufacturing sector from 2006 to 2020. Then, with simulated cost parameters, we examine the practical feasibility of identifying bilaterally implementable security cost-sharing arrangements in these alliances.

KEYWORDS
cooperative game, implementation theory, interdependent security, social responsibility, supply chain risk

1 | INTRODUCTION AND RELATED LITERATURE

Firms increasingly belong to a variety of interorganizational networks, such as complex supply chains, strategic alliances, or other types of partnerships. Membership in these networks can evidently yield economic benefits, but they also necessitate substantial additional security investments due to increased exposure to interdependent or contagion risks (Kunreuther & Heal, 2003). For instance, in January 2013, the European food industry endured a horse-meat contamination scandal (Lawrence, 2013). Meat products from several retailers and fast-food chains in the United Kingdom and Ireland, advertised as containing beef, were discovered upon testing to have been contaminated with horse-meat. Further investigation revealed that in the complex meat supply networks, with contractors and subcontractors spread all across Europe, a particular supplier had indulged in deliberate contamination in a bid to cut costs. Several retailers, including Britain's largest retailer, TESCO, that had sourced the contaminated meat, faced economic repercussions from a drop in sales and reputational harm. Other notable cases of supply contamination include the adulteration of milk with melamine (Levi et al., 2020; Mu et al., 2016) and the 2008 heparin adulteration scandal (Babich & Tang, 2012). Contamination in supply networks, upon discovery, typically

Accepted by Jayashankar Swaminathan, after two revisions.

results in product recalls, regulatory fines, and brand equity loss, often entailing substantial costs for the concerned firms.

Besides supply networks, interdependent risks can arise in other contexts too. For instance, businesses have a growing recognition that they bear a social responsibility to secure their consumer data from cyber threats (Pollach, 2011). Malware infecting the systems of a company in an interfirm network can gain access to the IT systems of its partner firms. Due to poor cyber-security practices by partner firms, companies such as Target and Home Depot have been the victims of high-profile data and privacy breaches (McAfee, 2015). In today's highly interconnected networks, risks like contamination in food supply chains or consumer data breaches assume an interdependent nature. That is, the risks faced by a firm depend not only on internal risks arising from their own operations but also on the risk transferred from partner firms in the network. Further, the above examples involve risks transferred between networked partners with ongoing and frequent repeated interactions. Thus, a firm vulnerable to internal risks is near-certain to transfer this risk to its partner firms if these partners do not take appropriate remedial actions.

Therefore, to secure themselves against interdependent risks, two general strategies are available to networked firms. First, firms in the network can choose to invest cooperatively in securing themselves, thereby removing sources of risk. Second, alternatively, firms can choose to independently secure themselves by eliminating risk from internal operations and then investing in security across the links that connect them to the other firms in the network. So, for example, firms could cooperatively share the costs of supplier quality improvements, thereby investing in suppliers' embracing responsible operational practices. Alternatively, a retailer can implement quality standards for internal processes and, simultaneously, inspect and quality test incoming products supplied by direct partners. The latter would correspond to the *independent security strategy*, while the former corresponds to the *cooperative security strategy*.

Security against interdependent risks is associated with positive externalities since other firms are benefited from the presence of a secured firm in the network. This would intuitively suggest that cooperative network-wide security against interdependent risks can be a cost-effective strategy as compared to each firm in the network independently securing itself. However, cooperation can be hindered by disagreements over cost-sharing arrangements. Firms, in general, are heterogeneous, both, in the costs they incur to secure themselves as well as in the penalties that they may face in case of a realized risk. Thus, a priori, it is not clear whether there will always exist a stable and fair sharing of security costs that can sustain network-wide cooperation. Furthermore, networked firms typically have visibility and mechanisms to cooperate and monitor with only immediate partners. For instance, extended multitier supply chains are often associated with a loss in visibility over firms further away in the network (Caro et al., 2021). Thus, it is also unclear whether one can find suitable mechanisms to implement cost-sharing arrangements that circumvent coordination across firms that are not immediate or direct partners.[1]

To address these issues, in this paper, we consider an interdependent security model on a network and an associated cost-sharing game. In our model, as motivated above, firms face an intrinsic risk from their internal operations and an extrinsic risk from their unsecured partners in the network. Firms in the network are heterogeneous in the costs they incur to secure themselves and the penalties they face in case of an actualized threat.

Further, we also consider our network security model under differing informational assumptions. In our *private information model*, we assume that all cost parameters are privately known to players. So, in the absence of explicit cooperation, each firm's security actions cannot be observed or inferred by other firms in the network. This private information assumption is a marked distinction from existing models of interdependent security in the literature, which typically assume that various model parameters and actions are public information. In several real-world contexts, in the absence of formal mechanisms for cooperation, firms are neither aware of the security efforts undertaken by other firms nor can they infer their efforts since the underlying cost structures are typically private information. However, in certain other scenarios, it would be more reasonable to assume that firms are indeed aware of the security costs of other firms in the network. Therefore, we also analyze our network security model with the alternative informational assumption wherein efforts and cost structures are *public information*. Further, studying these two extreme informational assumptions also permits us to separate the benefits of cooperation arising from interdependence and information acquisition. In the Supporting Information, we also consider a more general hybrid model, the *partial information model*, where, as in practice, due to regulatory requirements or strategic disclosures, the cost parameters and efforts of some firms are publicly known whereas the costs and efforts of other firms are only known privately.

The network-optimal security strategy under all informational assumptions is identical, and we demonstrate that it can be computed in polynomial time using a minimum weighted cut network-flow algorithm. Then, we adopt a cooperative game-theoretic approach to assess whether agents have an incentive to cooperate across the entire network and share the security investment costs. We show that, under the private information setting, agents have a clear incentive to cooperate globally, that is, form the grand coalition and share the resulting security costs. However, with even some information being public in the network, we show that, in general, there do not exist cost-sharing mechanisms that can ensure the stability of the grand coalition. This can be explained by two drivers: first, with public information, the benefit from additional information acquisition is lowered. Thus, the benefits from cooperative security in the public information setting are arguably lower. Second, public information engenders free-riding since firms can now anticipate and observe the security actions of other firms in the network and benefit from the cooperation of other firms in the network without participating in the grand coalition and sharing security costs. In similar cooperative settings with externalities, free-rider

concerns are acknowledged as a fundamental reason often precluding the stability of the grand coalition (see, e.g., Yi, 1997).

Importantly, we then introduce the notion of *bilateral implementability*. A cost-sharing arrangement is said to be bilaterally implementable if it can be enforced by a series of bilateral cost-sharing agreements between only direct partners in the network. Bilaterally implementable cost-sharing mechanisms are resistant to the aforementioned limitations of network visibility and control. It is generally assumed, for example, in managing supply chains that it is easier for firms to contract with their immediate suppliers with whom they share direct relationships and that it is more challenging to gain visibility, manage, and contract with deep-tier suppliers (see, e.g., L. Huang et al., 2020; Dong et al., 2022). We propose a novel security cost-sharing mechanism, the *agreeable allocation*, which is a restricted variant of the Shapley value allocation (Shapley, 1971). We then demonstrate that the agreeable allocation satisfies notions of stability, is formalizably fair, and unlike the Shapley value, is easily computable, and always bilaterally implementable. However, the agreeable allocation may not always exist. We then construct δ-agreeable allocations that satisfy a generalized notion of $(\delta+1)$-lateral implementability, for an integer $\delta \geq 1$, whereby firms that are at a distance of at most δ from each other in the network can enter into cost-sharing agreements. When $\delta = 1$, we recover bilateral implementability. This allows us to delineate a hierarchy of cost-sharing mechanisms such that as δ increases (i.e., firms that are farther away from each other in the network are allowed to cooperate), the corresponding δ-agreeable allocation is more likely to exist.

To analyze the effects of network structure on the existence of the agreeable allocation, we consider the special case of quasi-homogeneous networks, that is, networks where the security cost parameters are equal. We then provide a structural graph-theoretic characterization for the existence of the agreeable allocation in these networks. Specifically, we show that the local density of networks plays a key role in determining whether the agreeable allocation exists.

In summary, one can view our work in both descriptive and normative terms. Descriptively, we observe that network-wide security cooperation is efficient and, in some cases, this cooperation can be sustained with suitable cost-sharing arrangements. However, when concerns pertaining to computability and implementability of these cost-sharing mechanisms are incorporated, network-wide security cooperation is rendered more challenging. Normatively, via our analysis of the agreeable allocation and its extensions, we are able to provide insights into when and how these implementation challenges can be surmounted.

1.1 | Overview of related literature

This work is related to three distinct streams of literature. First, it contributes to extant work on social responsibility and risk management in supply chains. Second, our work is closely tied to interdependent security models introduced by Kunreuther and Heal (2003). One of our aims is to bridge these two bodies of literature. Finally, our work adds to the growing literature on applications of cooperative game theory to operations management.

1.1.1 | Supply chain social responsibility and risk management

There is a vast literature investigating the role of several instruments such as auditing (Caro et al., 2018; J. Chen et al., 2020; Fang & Cho, 2020; Plambeck & Taylor, 2016), inspection and testing (Babich & Tang, 2012; Lee & Li, 2018), and more recently, contracts (Dhingra & Krishnan, 2021), in mitigating social responsibility risks associated with extended global supply chains. We refer the interested reader to Dawande and Qi (2021) for a recent review. While previously, most of this literature dealt with two firm or dyadic scenarios, recently, several studies also deal with multitier supply chains, for example, supply networks with three tiers or other network structures (J. Chen et al., 2020; L. Huang et al., 2020; Zhang et al., 2022). Also closely related to our work, Feng et al. (2021) study the implementation of environmental and social responsibility (ESR) programs in general supply networks and gain sharing via a bilateral bargaining framework that generalizes a conventional Shapley value based cooperative-game theoretic approach. Recently, Blaettchen et al. (2021) also study the optimal adoption seeding of traceability technologies which carry several implications for sustainable practices in supply networks. While we view our work as contributing to this stream of literature, we note that it bears some differences. For instance, we consider a general network structure and do not impose any structural assumptions. Second, our work deals with only interdependent risks. That is risks that are contagion risks spreading via the network. These scenarios include cases such as food contamination risks or data breach threats as motivated in the introduction.

1.1.2 | Interdependent security

In terms of model development, our work is most closely related to the interdependent security literature. Interdependent security models were introduced by Kunreuther and Heal (2003) and have since spawned a rich literature on the intersection of economics and computer science that studies various related models (see, e.g., Laszka et al., 2014, for a review). In these models, as in ours, the security of agents depends on an agent's own actions (direct risk, or as we term it, intrinsic risk) and those of other agents (indirect or extrinsic risk). The present work aims to bridge the interdependent security literature with the rich stream of work on socially responsible operations in supply networks. While this research stream inspires our model, our work differs from existing literature in some crucial ways. First, in several of the existing models, the agents can only curb their own intrinsic risk and cannot mitigate extrinsic risks. Second, a majority

of the interdependent security literature adopts a noncooperative (game-theoretic) perspective. They assume that players in the network act to secure themselves independently and then characterize and compute the noncooperative equilibria of these games. Kearns and Ortiz (2003) and Chan et al. (2012) develop algorithms to compute the equilibria of classes of interdependent security games. Heal and Kunreuther (2007) also consider the Nash equilibria of such games and study conditions to tipping suboptimal equilibria to an optimal one. Chan and Ortiz (2014) consider a more general model where agents can influence the transfer of extrinsic risk and then analyze equilibria computations. However, this literature largely ignores issues of cooperation in networks and the problem of when and how cooperation can be sustained. In practice, agents can and indeed do cooperatively secure themselves against interdependent risks. This, therefore, is the central focus of this present paper.

1.1.3 | Cooperative game theory in operations management

Finally, we also contribute to the growing body of work dealing with applying cooperative game theory to problems in operations management. For a review of this literature, we refer the reader to Nagarajan and Sošić (2008). Benefits of cooperation can be realized and therefore studied in several diverse settings. Some recent applications include inventory pooling (Kemahlıoğlu-Ziya & Bartholdi, 2011), inventory transshipments (Granot & Sošić, 2003; Sošić, 2006), demand information sharing (Leng & Parlar, 2009), supplier alliances to mitigate order default risk (X. Huang et al., 2016), production schedule coordination (Aydinliyim & Vairaktarakis, 2010), supply chain emissions management and reduction (Gopalakrishnan et al., 2021a, 2021b), recycling (Gui et al., 2018; Tian et al., 2020), humanitarian operations (Ergun et al., 2014), vaccine distribution (Westerink-Duijzer et al., 2020), and so forth. Related to our work, Mu et al. (2019) study quality management in milk cooperatives. In dairy cooperatives, individual farmers can shirk on quality and free-ride on the higher quality milk produced by other farmers in the cooperative. Mu et al. (2019), therefore, develop a revenue allocation rule that achieves quantity and quality efficiency with minimal testing while incorporating other practical implementation considerations.

2 | A NETWORK SECURITY MODEL

We consider a set of heterogeneous players[2] denoted by N. Following standard graph-theoretic notation, let us suppose that the players occupy a network denoted as $\mathbb{G} = (N, A)$. The node set N of the network coincides with the set of players with each player occupying a unique corresponding node in $\mathbb{G}$. An arc $(i, j) \in A$ for $i, j \in N$ represents a directed link from the player i to the player j. The set of arcs in the network is denoted by A. Let $N^+(i)$ denote the set of players in N to which i is connected by an outgoing arc $(i, j) \in A$, and, similarly, let $N^-(i)$ be the set of players $j \in N$ such that the arc $(j, i) \in A$. Further, let $N(i) := N^+(i) \cup N^-(i)$.

Each player faces two independent sources of risk: an *intrinsic risk* from its own operations and an *extrinsic risk* transferred from its partnerships with unsecured players.[3] We assume the cost incurred by player i to secure itself against intrinsic risks is given by θ_i. Further, the cost incurred by i to secure itself against the extrinsic risk transferred from a partner in the network j is denoted by ξ_{ji}. Each player i exerts binary actions, $x_i \in \{0, 1\}$, and $y_{ji} \in \{0, 1\}$ for all $j \in N^-(i)$, corresponding to whether to secure itself against its own intrinsic risk and extrinsic risk from its partners, respectively. Since different players may face differing penalties (in regulatory fines or reputational damage) in the case of a realized risk, we assume an *unsecured* player i faces an expected penalty of L_i. A *secured* player faces a zero penalty. We will subsequently clarify when a player is said to be *secured* and *unsecured*, respectively.

As outlined in Section 1, firms can derive two distinct advantages from cooperative security in networks: first, the benefit of interdependence, which involves internalizing the positive externality of security, and, second, the advantage of information acquisition. Accordingly, we first consider two extreme informational assumptions, a *private information* model where each player, in the absence of cooperation, is aware of and can observe only its own security cost parameters and actions. At the other extreme, we also consider the more traditional informational assumption of *public information* where, even in the absence of cooperation, each player can observe the costs and actions of all other players in the network.

2.1 | Private information model

In the private information model, we assume that all cost parameters including the cost of securing against intrinsic risk, θ_i, and the expected penalty in case of a realized risk, L_i, are private information known only to player i. Similarly, the cost, ξ_{ji}, to secure the directed link between players j and i is assumed to be known only to players i and j. This private information assumption is a departure from several existing models of interdependent security. Specifically, the private information assumption implies that in the absence of explicit cooperation between players i and j, neither can observe or infer the actions of the other. Thus, in this scenario, we can formally define the *information set* of a player i acting independently as $I(i, \{i\}) = \{\theta_i, \xi_{ij}, \xi_{ji}, L_i, x_i, y_{ji} : j \in N^-(i)\}$. Therefore, in this scenario, the information set of player $i \in N$ who cooperates with the set of players $i \in S \subseteq N$ expands and is given by $I(i, S) = \cup_{j \in S} I(j, \{j\}) = \{\theta_j, \xi_{kj}, \xi_{jk}, L_j, x_j, y_{kj} : j \in S, k \in N^-(j)\}$.

2.2 | Public information model

In contrast, in the public information model, we assume that all firms can observe each other's cost parameters

and security actions even in the absence of cooperation. Then, the information set of a player i acting independently is $I(i,\{i\}) = \{\theta_j, \xi_{jk}, \xi_{kj}, L_j, x_j, y_{jk} : j \in N, k \in N^-(j)\}$. Therefore, in the public information scenario, $I(i,S) = I(i,\{i\})$, and firms upon cooperation do not derive any benefits from additional information acquisition. By analyzing and comparing these two extreme informational assumptions, we can comment on the benefits from cooperation along the two dimensions of interdependence and information acquisition.

2.3 | Partial information model

In practice, even in the absence of explicit cooperation, the security costs and actions of certain firms may be public knowledge, due to regulatory requirements or strategic disclosures, whereas the costs and actions of other firms may only be known privately. Thus, we also consider a more general partial information model which assumes that the costs and actions of a subset of firms, $\mathcal{P} \subseteq N$ are publicly known to all firms in the network whereas the costs and actions of firms in $N \backslash \mathcal{P}$ are only privately known. Therefore, in this scenario, $I(i,\{i\}) = \{\theta_j, \xi_{jk}, \xi_{kj}, L_j, x_j, y_{jk} : j \in \mathcal{P} \cup \{i\}, k \in N^-(j)\}$. This more general hybrid model subsumes both the private and public information models described above. Clearly, when $\mathcal{P} = \emptyset$ and $\mathcal{P} = N$, we recover the private and public information models, respectively. In the interest of expositional clarity and brevity, we consider the private information and public information models in the paper and extend the discussion to the general partial information model in the Supporting Information Section EC.4.

2.4 | Security actions

Players in the network choose security actions, $x_i \in \{0,1\}$, and $y_{ji} \in \{0,1\}$ for all $i \in N$ and $j \in N^-(i)$ after considering the relevant trade-off between the costs of security and the expected penalty in case of a realized risk. In order to do so, each player first forms beliefs on the *security states* of other firms in the network. That is, a player i, cooperating with players in S and with the information set $I(i,S)$, forms a belief on the security state of $j \in N$ denoted by $\sigma_{ji}(I(i,S)) \in \{0,1\}$, where $\sigma_{ji} = 0$ means player i believes j to be unsecured, and if $\sigma_{ji} = 1$ then i believes j is secured. We will subsequently clarify how players form beliefs on the security states of other firms in the network. Then, player i chooses security actions x_i and y_{ji} accordingly to determine its own security state based on its beliefs. Since interdependent risks are transferable across partners, a player i identifies itself as secured, that is, $\sigma_i = 1$, if and only if it is secured against its own intrinsic risk, that is, $x_i = 1$, and further, is also secured against extrinsic risks, that is, $y_{ji} = 1$ for all players $j \in N^-(i)$ who it believes to be unsecured. For clarity, we note that the security state σ_i of player i as a function of its own security actions, given its information set and its beliefs on the security states of its network partners, satisfies the following:

$$\sigma_i(x_i, \boldsymbol{y}_i | I(i,S)) = \begin{cases} 0, & \text{if } x_i \prod_{\substack{j \in N^-(i) \\ \sigma_{ji}=0}} y_{ji} = 0, \\ 1, & \text{otherwise.} \end{cases} \tag{1}$$

Thus, the expected security cost incurred by a player i is given as follows:

$$U_i(x_i, \boldsymbol{y}_i | I(i,S)) = L_i(1 - \sigma_i(x_i, \boldsymbol{y}_i | I(i,S))) + \theta_i x_i + \sum_{j \in N^-(i)} \xi_{ji} y_{ji}. \tag{2}$$

The first term in (2) corresponds to the expected penalty from a realized risk and is incurred only when the player i is unsecured. The second and third terms correspond to the costs of securing itself against intrinsic risks, and extrinsic risks from unsecured partners, respectively.

In Sections 3 and 4, we analyze cooperative security strategies and the associated security cost-sharing problem in the private information model whereas in Section 6 we study the public information model. This sequence is chosen for expositional clarity. Further, in the interest of parsimony, we relegate the analysis under the general *partial information* model where each player acting independently is aware of the cost parameters and actions for only a subset of players to the Supporting Information Section EC.4.

3 | SECURITY STRATEGIES UNDER PRIVATE INFORMATION

Under the private information assumption, since a player cannot observe or infer the security actions of other players, we assume a player i forms a worst-case belief on the security states of players it does not explicitly cooperate with. That is, a player i cooperating with the set of players $S \subseteq N$ forms the worst-case belief that $\sigma_{ji} = 0$ for all players $j \notin S$. Therefore, i identifies itself as secured if and only if it is secured against its own intrinsic risk, $x_i = 1$, and, further, is also secured against extrinsic risks, $y_{ji} = 1$ for all j such that $\sigma_{ji} = 0$, that is, (i) for j not in S and (ii) for j in S who are themselves not secured. Therefore, in the private information model, the security state of i is denoted by $\sigma_i \in \{0,1\}$, where $\sigma_i = 0$ means, in the worst case, player i is unsecured, and if $\sigma_i = 1$, then i is secured in the worst case. Similar worst-case considerations are commonly employed in diverse network security applications (see, e.g., a review on planning for supply network disruptions by Snyder et al., 2006).

We now consider two forms of security strategies in the network: the *independent security strategy* and the *network-optimal security strategy*. While the former corresponds to the no-cooperation, that is, individually rational scenario, the latter corresponds to the full cooperation, that is, the network-optimal situation. In Section 4, we will consider all intermediate cooperative security strategies, that is, where a subset of firms in the network cooperatively secure themselves.

3.1 | Independent security strategy

Since the players are not cooperating with each other on their security actions, as noted previously, the information set of each player $i \in N$, $I(i, \{i\})$, only contains its own actions, expected penalty, and security costs. Then, player i is said to be independently secured if U_i, as defined in (2), is minimized when $\sigma_i = 1$, for a suitable choice of x_i and $\boldsymbol{y}_i$. The set of all players in N which are independently secured is denoted by $S_{\mathcal{I}}$. The following proposition characterizes when a player is independently secured. All proofs are provided in the Supporting Information.

Proposition 1. *A player $i \in S_{\mathcal{I}}$ if and only if $L_i \geq \theta_i + \sum_{j \in N^-(i)} \xi_{ji}$. Further, then, $x_i = y_{ji} = 1$ for all $j \in N^-(i)$.*

The above proposition captures two straightforward notions in the private information setting: (i) the independent security strategy is based on a simple trade-off between the cost of security and the expected penalty incurred from not securing itself, (ii) for an agent acting independently, it is not optimal to partially invest in securing some links and not others.

3.2 | Network-optimal security strategy

In this setting of full network-wide cooperation, the information set of each player contains all the security costs and expected penalties of all other players in the network. The players act to minimize the total expected security cost of the network.

$$\begin{aligned} U(\mathbb{G}) &= \min_{\boldsymbol{x},\boldsymbol{y}} \quad \sum_{i \in N} U_i(x_i, \boldsymbol{y}_i | I(i,N)) \\ &= \min_{\boldsymbol{x},\boldsymbol{y}} \quad \sum_{i \in N} \Bigg(L_i(1 - \sigma_i(x_i, \boldsymbol{y}_i | I(i,N))) \\ &\quad + \theta_i x_i + \sum_{j \in N^-(i)} \xi_{ji} y_{ji} \Bigg). \end{aligned} \tag{3}$$

We denote the set of all players in N which are secured, that is, $\sigma_i = 1$, under the above network-optimal security strategy by $S_\star$. We first observe that all players that opt to be secured under the independent security strategy continue to be secured under the network-optimal strategy.

Proposition 2. *Every player independently secured is also secured under the network-optimal security strategy, $S_\star \supseteq S_{\mathcal{I}}$.*

However, the positive externalities, inherent to this context, may result in certain nodes being secured under the network-optimal security strategy which are unsecured when acting independently. That is, we note that the above inclusion can be strict. We demonstrate this with Example EC.1 in the Supporting Information.

We now provide a key result demonstrating that the network-optimal security strategy and equivalently, $U(\mathbb{G})$, can be computed via a network-flow algorithm. The algorithm relies on the construction of an auxiliary directed network $\mathbb{G}^*$. We then establish a connection between the network-optimal security strategy in $\mathbb{G}$ and the minimum weight s-ℓ cut problem in $\mathbb{G}^*$.

3.3 | Construction of the auxiliary network $\mathbb{G}^*$

The node set of $\mathbb{G}^*$ is given by $N \cup \{s, \ell\}$, where s and ℓ are two additional nodes not present in the original network $\mathbb{G}$. The nodes s and ℓ represent the source and sink of the network $\mathbb{G}^*$, respectively. The arc set of $\mathbb{G}^*$ consists of (i) arcs from s to each node $i \in N$ with weights θ_i, (ii) arcs from $i \in N$ to $j \in N^+(i)$ with weights ξ_{ij}, and (iii) arcs from $i \in N$ to ℓ with weights L_i. The construction of the auxiliary network is illustrated in Figure 1.

Theorem 1. *Suppose the minimum weight cut separating s and ℓ partitions the nodes of $\mathbb{G}^*$ into X and $\overline{X}$ such that $s \in X$. Then $S_\star = N \setminus X$. Further, $U(\mathbb{G})$ is the weight of the cut $(X, \overline{X})$.*

Also, from (1), it follows that if $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$ denote the network-optimal security actions of the players, then $x_i^* = 1$ if and only if $i \in \overline{X}$ and $y_{ji} = 1$ if and only if $i \in X$, $j \in \overline{X}$. Therefore, from Theorem 1, we also immediately obtain the network-optimal security strategy. Now, note that the directed network $\mathbb{G}^*$ has $O(|N|)$ nodes and $O(|N| + |A|)$ arcs. Thus, from the push-relabel-algorithm (Goldberg & Tarjan, 1988), we immediately obtain the following corollary.

Corollary 1. *$S_\star$ can be computed in $O((n^2 + mn)\log(n/m + n))$ time, where $n = |N|$ and $m = |A|$.*

In the private information model, the network-optimal security strategy resolves two distinct kinds of inefficiencies engendered by the individually rational security strategies of the players. The first inefficiency arises from the canonical underinvestment of efforts resulting from a failure to internalize positive externalities. This is well recognized in the interdependent security literature (see, e.g., Acemoglu et al., 2016). Therefore, some agents for whom it was individually rational to not invest in security efforts are now secured since these erstwhile externalities are now internalized in the network-level optimization. This reflects the strategic complementarity inherent in situations with interdependent risks. The second source of inefficiency arises, in the private information model, as a consequence of security costs being privately held information. Equivalently, the noninferability of security efforts of a player by other players who are not cooperating with it results in the inefficient duplication of security investments across the network. This provides an economic rationale for anecdotal evidence from diverse

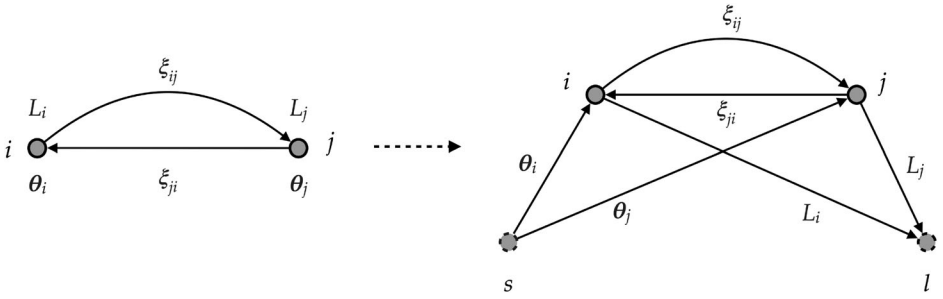

FIGURE 1 Auxiliary network $\mathbb{G}^*$.

supply chain security contexts that bear out this source of inefficiency (ASEM, 2013).

Finally, we note the necessity of cost-sharing mechanisms in order to implement the network-optimal security strategy. For a player in the network, given the security states of all of its direct partner firms, the network-optimal security action is not necessarily individually rational. That is, the network-optimal security strategy is not always a Nash equilibrium strategy as demonstrated by Example EC.2 provided in the Supporting Information.

4 | SECURITY COST SHARING MECHANISMS

The next natural question is therefore to ask whether network-wide security cooperation in the private information model can be sustained with suitable cost-sharing mechanisms. Equivalently, we are interested in finding whether and when cooperation can be made individually rational and the network-wide efficiency gains can be shared among the firms in a *stable* and *fair* manner. The field of cooperative game theory is well suited to address these questions. Towards that end, we first briefly review some cooperative game theory preliminaries.

Cooperative game theory primarily addresses the question of whether cooperation can be sustained across a group of agents and, closely tied to this, is the problem of *fairly* sharing or allocation of profits (or cost savings) obtained via cooperation between those agents. A cooperative game is defined by (N, c), where N is the set of players in the game and $c(\cdot)$ is a characteristic function that associates with every subset (or, coalition) $S \subseteq N$ a corresponding cost $c(S)$. The subset consisting of all players, that is, the set N itself is known as the grand coalition. An object of frequent interest is whether the grand coalition will form and whether it remains rational for individual players, or groups of players, to remain in the grand coalition. In this work, we will only deal with cost games, that is, where $c(S)$ is the cost incurred by coalition S, and players act to minimize their costs. A cooperative game (N, c) is said to be *subadditive* if the characteristic function satisfies $c(S) + c(T) \geq c(S \cup T)$ for $S, T \subseteq N$. Subadditivity can loosely be interpreted as offering an incentive for disjoint coalitions to cooperate. Another important property that a cooperative game can satisfy is convexity. The convexity property is stronger than the subadditivity property, and it loosely captures the intuition that as a coalition grows larger, the greater the incentive for other players to join it. Formally, $c(S) + c(T) \geq c(S \cup T) + c(S \cap T)$ for $S, T \subseteq N$.

4.1 | Interdependent security cost sharing

Consider the set of agents N situated on the graph $\mathbb{G}$. Previously, the two security strategies considered represented the two extremes corresponding to no-cooperation and full-cooperation settings. We now extend the discussion to consider all intermediate levels of cooperation. That is, for any subset of agents, $S \subseteq N$, we define the *coalition-optimal security strategy* as that which minimizes the security cost of a cooperating set of agents S,

$$\begin{aligned} c(S) &= \min_{\boldsymbol{x},\boldsymbol{y}} \quad \sum_{i \in S} U_i(x_i, \boldsymbol{y}_i | I(i,S)) \\ &= \min_{\boldsymbol{x},\boldsymbol{y}} \quad \sum_{i \in S} \left(L_i(1 - \sigma_i(x_i, \boldsymbol{y}_i | I(i,S)) + \theta_i x_i + \sum_{j \in N^-(i)} \xi_{ji} y_{ji} \right). \end{aligned} \tag{4}$$

We define an indicator function Y_S^i for player i belonging to a coalition S that indicates whether player i is secured under the coalition-optimal security strategy for S in the private information model. Formally, $\mathrm{Y}_S^i = \sigma_i(\widetilde{x}_i, \widetilde{\boldsymbol{y}}_i | I(i,S))$, where $\widetilde{x}_i$ and $\widetilde{\boldsymbol{y}}_i$ denote the optimal solutions to (4). Further, denote the set of players secured in S under the coalition-optimal security strategy by $\mathrm{Y}(S)$. That is, $i \in \mathrm{Y}(S)$ if and only if $\mathrm{Y}_S^i = 1$. Clearly, $S \backslash \mathrm{Y}(S)$ are the players in S that are not secured under the coalition-optimal security strategy. Further, for clarity, note that $\mathrm{Y}(N) = S_\star$. The following result demonstrates a monotonicity property satisfied by the coalition-optimal security strategy that generalizes Proposition 2.

Proposition 3. *A player $i \in S$ that is secured under the coalition-optimal security strategy for a coalition $S \subseteq N$ is also secured under the coalition-optimal security strategy for a coalition $T \supseteq S$, that is, if $\mathrm{Y}_S^i = 1$, then $\mathrm{Y}_T^i = 1$.*

Further, the pair (N, c) defines a cooperative game which we term as the *interdependent security cost-sharing game*. This cost-sharing game corresponds to our network model

based on the private information assumption as clarified in Section 2. In Section 6, we will accordingly define and analyze the appropriate cost-sharing game for the public information setting.

The following proposition indicates that $c(S)$ can also be computed in polynomial time via a similar transformation to a minimum weight cut problem on the auxiliary graph $\mathbb{G}^*$ as in Theorem 1.

Proposition 4. *$c(S)$ is the weight of the minimum cut separating the node set $N\backslash S$ and the node ℓ in the auxiliary directed graph $\mathbb{G}^*$ and thus can be computed in polynomial time.*

An efficient security cost-sharing mechanism is defined as $\boldsymbol{\phi} : (N, c) \to \mathbb{R}^n$ such that $\sum_{i\in N} \phi_i = c(N)$. An efficient security cost-sharing mechanism is said to be a core allocation, that is, it belongs to the core if and only if it is rational for all subsets of players in N to remain in the grand coalition rather than deviate to form a coalition among themselves. That is, $\boldsymbol{\phi}$ is a core allocation if and only if, $\sum_{i\in S} \phi_i \leq c(S)\ \forall\ S \subseteq N$. The core of some cooperative games may be empty. An empty core will preclude the existence of *stable* cost-sharing arrangements. However, in cooperative games that are also convex, it is well known that the core of such games is nonempty (Shapley, 1971). The following theorem demonstrating the convexity of the interdependent security cost-sharing game, therefore, assumes significance since it guarantees the existence of a stable cost-sharing mechanism.

Theorem 2. *The coalition-optimal security cost, $c(S)$, is submodular in S. Thus, the interdependent security cost-sharing game (N, c) always admits a stable security cost-sharing mechanism.*

Before we proceed to derive and analyze specific security cost-sharing mechanisms, we observe that if a player is unsecured under the network-optimal security strategy, then, the player is allocated L_i by all stable cost-sharing arrangements as formally demonstrated in Lemma EC.1. Further, we also show that there exists a simple transformation of a network $\mathbb{G}$ where some players are unsecured under the network-optimal security strategy to another network $\mathbb{G}'$ where all players are secured in the network-optimal strategy and, further, there exists a one-to-one correspondence between the core allocations of the interdependent security games on $\mathbb{G}$ and $\mathbb{G}'$. Thus, Lemma EC.1 allows us to restrict our attention to networks $\mathbb{G}$ and associated cost parameter vectors such that all firms are secured under the network-optimal security strategy.

4.1.1 | Shapley value based security cost sharing

The convexity of (N, c) guarantees that a well-known and commonly employed allocation in cooperative games, the Shapley value (Shapley, 1953), belongs to the core. Beyond its membership in the core, the Shapley value also uniquely satisfies several natural fairness properties and has an axiomatic basis in general cooperative games. Formally, the Shapley value, Φ, allocates to a player i in a general cooperative game (N, c),

$$\Phi_i = \sum_{\{S:i\in S\}} \frac{(|S|-1)!(n-|S|)!}{n!}\big(c(S) - c(S\setminus\{i\})\big). \quad (5)$$

The Shapley value rewards players for their marginal contributions to various coalitions, and, to that extent, it can be argued as exemplifying a certain notion of fairness. Further, Φ is the unique *efficient* allocation characterized by the following properties (or axioms):

(i) *Symmetry property*: For players i and j such that for all subsets $S \subset N$, $i, j \notin S$, if $c(S\cup\{i\}) - c(S) = c(S\cup\{j\}) - c(S)$, then $\Phi_i = \Phi_j$.
(ii) *Null player property*: For player i such that $c(S\cup\{i\}) = c(S)$ for all $S \subset N$, then $\Phi_i = 0$.
(iii) *Additivity property*: The Shapley value, $\Phi^{1,2}$, of a cooperative game, $(N, c^1 + c^2)$, that is the sum of two cooperative games, (N, c^1) and (N, c^2), equals the sum of the Shapley values of the two games, Φ^1 and Φ^2, respectively.

Of these properties, we note that the symmetry property formalizes the idea that players which are "identical" in terms of their marginal contributions should receive an identical share of the value created by cooperation. This is, arguably, an innocent fairness criterion which, along with the marginal contribution interpretation discussed before, we shall return to later on in this work. The Shapley value is widely adopted as a cost sharing or a profit sharing, as the case may be, allocation method in diverse contexts, including several mentioned in Section 1.1, such as inventory pooling (Kemahlıoğlu-Ziya & Bartholdi, 2011), capacity allocation and scheduling (Aydinliyim & Vairaktarakis, 2010), group purchasing (R. R. Chen & Yin, 2010), disaster preparedness (Rodríguez-Pereira et al., 2021), and so forth. However, for our game, we establish a link between the computation of the Shapley value and the classical subset sum problem. In fact, this connection demonstrates that computing the Shapley value of interdependent security games is a computationally hard problem.

Theorem 3. *There is no polynomial time algorithm that computes the Shapley value for a given player in the interdependent security cost-sharing game unless $P = NP$.*

Further, from the proof of Theorem 3, we note that even for simple structures such as the assembly supply network, computing the Shapley value is hard. Beyond computational interest, the above result on the complexity of the Shapley value is of interest to us for reasons of implementation. In general, equilibrium concepts in noncooperative game theory

or solution concepts in cooperative games that are computationally intractable raise the question of feasibility of whether self-interested agents can identify and implement these mechanisms in practice.[4]

For a notable special case, however, the Shapley value can be computed easily. In fact, when the expected penalties, in case of a realized risk, are sufficiently large for all players, then the Shapley value has a straightforward closed form expression.

Theorem 4. *If $L_i > \theta_i + \sum_{j\in N^-(i)} \xi_{ji}$ for all $i \in N$, that is, if $S_{\mathcal{I}} = S_\star = N$, then, the Shapley value based security cost allocation to player $i \in N$ is given by*

$$\Phi_i = \theta_i + \sum_{j\in N^-(i)} \frac{\xi_{ji}}{2} - \sum_{j\in N^+(i)} \frac{\xi_{ij}}{2}. \tag{6}$$

In this scenario, when the expected penalties are sufficiently large, it is individually rational for all players to secure themselves (i.e., under the independent security strategy). That is, since all players choose to secure themselves even without cooperation, the network-optimal security strategy resolves only one kind of inefficiency, that arising from duplication of security efforts. Under the Shapley value based security cost-sharing mechanism, in this scenario, the cost savings from avoiding duplication of security efforts across each link are equally shared by both parties.

4.1.2 | Extreme core allocations

However, this still leaves open the question of whether, in general interfirm networks, there exist stable security cost-sharing arrangements sustaining network-wide cooperation that can also be computed easily. We now provide an affirmative answer to this question. Consider an arbitrary permutation π of the players in N. Then, we can define a cost-sharing allocation, x_π, corresponding to a permutation π as follows, $x_{\pi_i} = c(\{\pi_1, \pi_2, \dots, \pi_i\}) - c(\{\pi_1, \pi_2, \dots, \pi_{i-1}\}) \quad \forall i \in N$.

Proposition 5. *For every permutation π of N, the allocation x_π is an extreme point of the core of the interdependent security cost-sharing game and can be computed in polynomial time.*

The proof of Proposition 5 relies on the convexity of the game and the characterization of the core of convex games as developed by Shapley (1971). Further, we demonstrate that the extreme core points of the interdependent security cost-sharing game can be computed in polynomial time, thereby, allowing us to identify easily computable and stable security cost-sharing arrangements. However, it can easily be seen that extreme core allocations as identified in Proposition 5 do not satisfy a basic notion of fairness as embodied in the symmetry property introduced earlier.

Proposition 6. *The security cost-sharing allocation x_π does not satisfy the symmetry property.*

Our discussion, thus far, uncovers what appears to be an "impossible" trilemma: *stability*, *fairness*, and *implementability*. That is, when we simultaneously require a security cost-sharing arrangement to be stable (i.e., it must be individually and coalitionally rational), fair (in terms of a basic symmetry property), and implementable (in terms of ease of computability), it already proves to be too restrictive. Descriptively, this suggests why, although the welfare gains achieved by network-wide security cooperation can, in principle, be stably shared, we may still not observe such cooperation in practice. In the next section, we will delve deeper into implementability concerns. Further, and importantly, we will also attempt to find a satisfactory reconciliation of the divergence between stability, fairness, and implementability.

5 | BILATERAL AND MULTILATERAL IMPLEMENTABILITY

In Section 4, we considered a narrow version of implementability. Specifically, we presumed a security cost-sharing mechanism that is easily computable is implementable. However, implementing cost-sharing mechanisms via transfer payments across the network, even between firms that are not direct partners, is administratively challenging, perhaps even infeasible. Firms often have limited visibility let alone an ability to enter into cost-sharing arrangements with indirect network members. Therefore, in this section, we are prompted to study whether there exist stable and fair cost-sharing mechanisms that can be implemented via transfer payments only involving firms that are direct partners in the network. Indeed, since alliance networks are often comprised of a series of bilateral alliances in the first place, we develop a realistic bilateral implementation framework that can allow firms to sustain network-wide security cooperation against interdependent risks.[5]

To this end, we define the *bilateral implementability* of a cost-sharing allocation as follows. A cost-sharing allocation Ψ is bilaterally implementable if and only if for a given network $\mathbb{G}$ and associated cost parameter vectors $\{\boldsymbol{L}, \boldsymbol{\theta}, \boldsymbol{\xi}\}$, there exist differentiable functions $\{g_{ij} : j \in N(i)\}$ for each player $i \in N$ such that

$$\Psi_i = \sum_{j\in N(i)} g_{ij}(\theta_i, \theta_j, L_i, L_j, \xi_{ij}, \xi_{ji}) \tag{7}$$

for cost parameters belonging to an open ball $\mathcal{B}^\epsilon$ centered at $(\boldsymbol{L}, \boldsymbol{\theta}, \boldsymbol{\xi})$ of radius ϵ for some $\epsilon > 0$. That is, qualitatively, the security cost apportioned to each player i can be supported via verifiable transfer payments between only direct partners in the network. As discussed before, bilateral implementability obviates the need for transfer payments between firms not

direct partners in the network. And, consequently, since typically alliance networks expand via bilateral alliances, it also allows for sustaining network-wide cooperative security as the network structure evolves.

First, we examine the bilateral implementability of the Shapley value based security cost-sharing allocation discussed in Section 4. We introduce some definitions. For a given player $i \in N$, a set of players $P \subseteq N$ is said to be a *coalitionally rational security set* for i if i is secured in the coalitional optimal security strategy for the coalition $P \cup \{i\}$, that is, $i \in \mathrm{Y}(P \cup \{i\})$. We denote the set of all minimal[6] coalitionally rational security sets for player i by $\mathcal{G}(i)$ and further, $\overline{\mathcal{G}}(i) = \bigcup_{P \in \mathcal{G}(i)} P$.

Theorem 5. *Consider the Shapley value based security cost-sharing allocation* Φ.

(*i*) Φ *is bilaterally implementable if for all players* $i \in N$, $i \notin \overline{\mathcal{G}}(j)$ *for all* $j \in N(i)$ *such that* $|N(j)| > 1$.

(*ii*) Φ *is not bilaterally implementable if there exists a player* $i \in N$ *such that* $i \in \overline{\mathcal{G}}(j)$ *for some* $j \in N(i)$ *such that* $|N(j) \backslash N(i)| > 1$.

Theorem 5 provides characterizing conditions for when the Shapley value based cost-sharing arrangement is bilaterally implementable. Observe that minimal coalitionally rational security sets formalize the externalities that secured players induce on other players in the network. Therefore, roughly speaking, the above theorem demonstrates that as the extent of positive externalities of security in the network increases, the Shapley value based security cost-sharing fails to be bilaterally implementable. As a corollary, we observe that for the special case discussed in Theorem 4, the Shapley value cost-sharing mechanism is clearly bilaterally implementable.

Theorem 5, in conjunction with Theorem 3, arguably also demonstrates the impracticality of adopting a Shapley value based security cost-sharing arrangement in all but a narrow class of networks. Specifically, since it is neither computable efficiently nor bilaterally implementable, in general, we argue that this renders it contextually untenable. We now propose a novel security cost-sharing mechanism that builds on the extreme core allocations considered in Proposition 5.

5.1 | Extreme core allocations and the agreeable allocation

In light of Lemma EC.1, we limit our attention to networks where all firms are secured in the grand coalition. We further recall the previously defined indicator function Y_S^i for player $i \in S$ that indicates whether player i is secured under the coalition-optimal security strategy for S. That is, $\mathrm{Y}_S^i = \sigma_i(\widetilde{x}_i, \widetilde{\mathbf{y}}_i | I(i,S))$, where $\widetilde{x}_i$ and $\widetilde{\mathbf{y}}_i$ denote the optimal solutions to (4). We now recursively define a finite family of mutually exclusive sets $\mathcal{S} = \{\mathcal{S}_1, \ldots, \mathcal{S}_\ell\}$ of players in the network, where $\mathcal{S}_1 = \{i \in N : \mathrm{Y}_{\{i\}}^i = 1\} = S_I$. For $k > 1$, we define $\mathcal{S}_k$ recursively as

$$\mathcal{S}_k = \left\{ i \in N \setminus \overline{\mathcal{S}_{k-1}} : \mathrm{Y}^i_{\overline{\mathcal{S}_{k-1}} \cup \{i\}} = 1 \right\}, \tag{8}$$

where $\overline{\mathcal{S}_{k-1}} = \mathcal{S}_1 \cup \cdots \cup \mathcal{S}_{k-1}$. In other words, $\mathcal{S}_1$ contains the players that are secured even under the independent security strategy, that is, it is optimal for these players to secure themselves even when operating independently. Further, $\mathcal{S}_2$ contains players that will be secured conditional on being in a coalition with players in $\mathcal{S}_1$, and so forth. Also note that if $\mathcal{S}_k$ is a null set, then, so is $\mathcal{S}_{k+1}$. Suppose there exists $\ell \in \mathbb{Z}$ such that $\overline{\mathcal{S}_\ell} = N$, then the recursive procedure generating the family of sets terminates. Denote $s_k = |\overline{\mathcal{S}_k}|$ for $k = 1, \ldots, \ell$. Then, any permutation π of the players in N such that $\pi_1, \ldots, \pi_{s_1}$ is a permutation of players in $\mathcal{S}_1$, $\pi_{s_1+1}, \ldots, \pi_{s_2}$ is a permutation of players in $\mathcal{S}_2$, and so on up to, $\pi_{s_{\ell-1}+1}, \ldots, \pi_{s_\ell}$ is a permutation of players in $\mathcal{S}_\ell$ is defined as an *agreeable permutation*.

We note that it is possible in certain networks and associated cost parameter vectors for no $\ell \in \mathbb{Z}$ to exist such that $\overline{\mathcal{S}_\ell} = N$. In these cases, consequently, no agreeable permutation of the players in N will exist either. Nevertheless, when the players in N can be partitioned into the family of sets as described above, or equivalently, when an agreeable permutation of the players exists, we can demonstrate, as will be shown during the course of proving Theorem 6, that the extreme core allocation x_π corresponding to each agreeable permutation π of N is bilaterally implementable.

Furthermore, recall that extreme core allocations are not symmetric, therefore, arguably, violating a basic notion of fairness. To remedy this, we are now in a position to propose our novel security cost-sharing mechanism, the *agreeable allocation*, that is defined as the average of those extreme core allocations induced by all agreeable permutations of N.

Theorem 6. *The agreeable allocation of network-wide security costs, when it exists,* (*i*) *belongs to the core and is* (*ii*) *polynomial-time computable,* (*iii*) *symmetric, and* (*iv*) *bilaterally implementable. Further, it also satisfies* (*v*) *marginality and the* (*vi*) *null player property. Moreover, the security cost allocated to player* i *by the agreeable allocation* x^* *is given by*

$$x_i^\star = \theta_i + \sum_{\substack{j \in N^-(i) \\ j \in (N \setminus \overline{\mathcal{S}_k})}} \xi_{ji} - \sum_{\substack{j \in N^+(i) \\ j \in \overline{\mathcal{S}_{k-1}}}} \xi_{ij} + \sum_{\substack{j \in N^-(i) \\ j \in \mathcal{S}_k}} \frac{\xi_{ji}}{2} - \sum_{\substack{j \in N^+(i) \\ j \in \mathcal{S}_k}} \frac{\xi_{ij}}{2} \quad \text{for } i \in \mathcal{S}_k.$$

Observe that the network-wide security cost apportioned to each player by the agreeable allocation depends only on its own security cost parameters and that of its direct

partners, and, therefore, it is bilaterally implementable. Also, importantly, we note that the agreeable allocation attempts to resolve the tension between *stability*, *fairness*, and *implementability*. Since, it belongs to the core, when it exists, it is a stable allocation of security costs. Further, in contrast to extreme core allocations, since it satisfies symmetry and marginality, it is in accordance with basic axiomatic descriptions of fairness. Further, in contrast to the Shapley value based cost-sharing arrangement, since the agreeable allocation is computable in polynomial time, and, saliently, is bilaterally implementable, it also fares well with respect to implementability concerns. Finally, the closed-form expression for the agreeable allocation provided above allows for transparency in the manner in which it allocates the network-wide security costs to each individual firm. In fact, the algorithm to compute the agreeable allocation and the closed-form expression lend themselves naturally to a straight-forward implementation mechanism.

We also remark that for the special case considered in Theorem 4, that is, when $S_{\mathcal{I}} = S_{\star}$, the agreeable allocation exists and coincides with the Shapley value.

5.2 | Multilateral implementability and δ-agreeable allocations

The agreeable allocation is indeed appealing since its bilateral implementability minimizes the coordination challenges involved in sustaining the network-optimal security strategy. However, sometimes firms that are not direct partners may regardless cooperate via suitable transfer payments when it can be mutually beneficial. Consider a network $\mathbb{G}$ with associated cost parameter vectors $\{\boldsymbol{L}, \theta, \boldsymbol{\xi}\}$. Formally, for an integer $\delta \geq 1$, a cost-sharing allocation Ψ is said to be $(\delta+1)$-*laterally implementable* if and only if for cost parameters belonging to an open ball B^{ϵ} centered at $(\boldsymbol{L}, \theta, \boldsymbol{\xi})$ of radius ϵ for some $\epsilon > 0$, there exist differentiable functions $\{g_{ij} : j \in N, d(i,j) \leq \delta\}$ for each player $i \in N$ such that $\Psi_i = \Sigma_{j \in N, d(i,j) \leq \delta}\ g_{ij}$, where g_{ij} is a function solely of the security cost parameters of players i and j, and where $d(i,j)$ denotes the distance between nodes i and j in the network $\mathbb{G}$. That is, $(\delta+1)$-lateral implementability of a cost-sharing allocation permits transfer payments between players that are at a distance of at most δ in the network. As δ increases, we expect the coordination challenges associated with the cost-sharing mechanism to also increase.

While our general approach to construct a $(\delta+1)$-laterally implementable allocation bears some resemblance to the previous development of the agreeable allocation, there are substantial technical differences. In the interest of brevity, we provide these details in the Supporting Information, Section EC.2. Broadly, we first identify a subset of permutations of the players in N denoted as δ-agreeable permutations (Algorithm 2). A δ-agreeable permutation can be computed via a fixed parameter tractable algorithm with respect to δ (i.e., polynomial time in $|N|$ but not in δ). We then demonstrate that the extreme core allocations corresponding to each δ-agreeable permutation are $(\delta+1)$-laterally implementable (Proposition EC.1). We then define the δ-agreeable allocation as the average of extreme core allocations induced by all δ-agreeable permutations of N.

Theorem 7. *For a given integer $\delta \geq 1$, the δ-agreeable allocation, when it exists (i) belongs to the core,(ii) is symmetric, and is (iii) (δ+1)-laterally implementable. Further, it also satisfies (iv) marginality and the (v) null player property.*

The δ-agreeable allocation satisfies the generalized notion of $(\delta+1)$-lateral implementability while retaining the fairness and stability properties of the agreeable allocation. Since the number of δ-agreeable permutations can be exponential in $|N|$, the δ-agreeable allocation is, in general, not computable in polynomial time for $\delta > 1$. However, as noted above, the δ-agreeable allocation can be computed via a fixed parameter tractable algorithm, that is, polynomial time in $|N|$ for a given δ. In comparison, we note that the Shapley value allocation is also not, in general, computable in polynomial time but since it involves the consideration of all permutations of N unlike the δ-agreeable allocation which only considers a subset of permutations of players in N, the δ-agreeable allocation is, in comparison, computationally less expensive, especially so when $|N|$ is large and δ is a fixed small number. In Section EC.2, we also provide Example EC.3 that clarifies the computation of the δ-agreeable allocation and illustrates the notion of $(\delta+1)$-lateral implementability.

Theorem 8. *Consider the interdependent security cost-sharing game under private information.*

(*i*) *If for an integer $\delta \geq 1$, the δ-agreeable allocation exists, then the $(\delta+1)$-agreeable allocation also exists and coincides with the δ-agreeable allocation.*
(*ii*) *For every integer $\delta \geq 1$, there exist networks $\mathbb{G}$ with corresponding security cost parameters such that the δ-agreeable allocation does not exist but the $(\delta+1)$-agreeable allocation exists.*
(*iii*) *The n-agreeable allocation always exists where $n = |N|$.*
(*iv*) *The n-agreeable allocation coincides with the Shapley value allocation if and only if none of the δ-agreeable allocations exist for $\delta < n$.*

Theorem 8 clarifies a hierarchy of existence for δ-agreeable allocations. As δ increases, and firms that are farther away from each other in the network are allowed to cooperate with each other via suitable transfer payments, the δ-agreeable allocation is more likely to exist. However, naturally, as δ increases, arguably, the δ-agreeable allocation becomes more challenging to implement than the agreeable allocation since it requires coordination between firms that are farther away in the network. Further, it follows from Theorem 8(iv), and since, in general, the Shapley value allocation involves transfer payments between any two firms in the network, δ-agreeable allocations are (weakly) less challenging to implement than the Shapley value.

6 | NETWORK SECURITY MODEL WITH PUBLIC INFORMATION

In this section, we consider the public information model, as presented in Section 2, wherein all network cost parameters and actions are known to all players in the network. That is, the information set of every player i in any coalition $S \subseteq N$ includes the security cost parameters and actions of all players in the network, $I(i,S) = \{\theta_j, \xi_{kj}, L_j, x_j, y_{kj} : j \in N, k \in N^-(j)\}$. Further, since a player can observe and infer the security actions of all other players in the network, player i no longer needs to form a worst-case belief[7] on the security state of other players $j \in N$, that is, $\sigma_{ji} = \sigma_j$. And thus, firm i ends up minimizing its expected cost rather than its worst-case expected cost.

Characterizing the security strategy of a coalition, or even the independent security strategy, in the public information model poses some challenges. In our network security model, as is often the case in network games with public information (Galeotti et al., 2010), there could be multiple Nash equilibria. Further, in the public information setting, the actions of a player or a coalition also depend on the actions of other players, and, therefore, naturally on whether other players in the network are cooperating with each other. Therefore, we cannot analyze the security actions of a player or a coalition in isolation. We instead need to consider the cooperation structure across the entire network. This is in contrast to the interdependent security cost-sharing game developed in Section 3, wherein the security cost of a coalition S could be expressed independent of considering the actions of other players. Therefore, the interdependent security cost-sharing problem under public information is modeled as a *cooperative game in partition function form* (see, e.g., Fang & Cho, 2020; Hafalir, 2007). Formally, given a partition ρ of the players into disjoint coalitions whose union is N, the total security cost incurred by a coalition $S \in \rho$ in equilibrium is denoted by $\widehat{c}(S;\rho)$.

Again, we first consider the security actions of players when they are all acting independently. That is, ρ consists of singleton sets of players. Each player $i \in N$ considers its security actions independently but knows all cost parameters in the network and can therefore infer the security actions of other players. Let $\widehat{\Upsilon}^i_{\{i\};\rho}$ be an indicator function denoting the equilibrium security state of player i acting independently where ρ is the coalition structure with all players in independent singleton coalitions. To address the multiplicity of equilibrium outcomes, we adopt a specific equilibrium selection procedure. Initially, all players choose their security actions independently without regard to the actions of other players in the network. Then, in subsequent rounds, players reassess their actions given the actions of others in preceding rounds. This procedure[8] is formally described (Algorithm 3) in Supporting Information Section EC.3.

Algorithm 3 computes an equilibrium security state of player i, $\widehat{\Upsilon}^i_{\{i\};\rho}$, in polynomial time. Given a general coalition structure ρ, we denote an equilibrium security state of player i in coalition S by $\widehat{\Upsilon}^i_{S;\rho}$. The equilibrium selection procedure described above for the case of independent coalitions can similarly be extended (Algorithm 4) to compute, in polynomial time, an equilibrium security strategy for a coalition $S \subseteq N$ with a general partition ρ of N with $S \in \rho$.

We then obtain the total security cost of a coalition S belonging to a general coalition structure ρ of N, $\widehat{c}(S;\rho)$, as follows:

$$\widehat{c}(S;\rho) = \sum_{i \in S} \left(L_i(1 - \widehat{\Upsilon}^i_{S;\rho}) + \theta_i \widehat{\Upsilon}^i_{S;\rho} + \sum_{\substack{(j,i) \in A \\ \widehat{\Upsilon}^i_{S;\rho}=1, \widehat{\Upsilon}^j_{T;\rho}=0}} \xi_{ji} \right), \tag{9}$$

where S and T are (possibly identical) coalitions in ρ with $i \in S$ and $j \in T$. For clarity, we note that for the grand coalition structure ρ^*, that is, when all players cooperate with each other, the total security cost under the public information and private information settings is equal, $\widehat{c}(N;\rho^*) = c(N)$. This is since even under the private information setting all players in the grand coalition are aware of all security cost parameters in the network.

We demonstrate that in the interdependent security cost-sharing game under public information, $(N,\widehat{c})$, the grand coalition is not necessarily stable. This is in contrast to our earlier result (Theorem 2) that there always exists a stable security cost-sharing mechanism under the private information setting. This can be explained by two drivers. First, in the public information setting, one of the benefits of cooperative security, the benefit from additional information acquisition is removed. Thus, the benefits from cooperative security in the public information setting are arguably lower. Second, public information engenders free-riding since firms can now anticipate and observe the security actions of other firms in the network and benefit from the cooperation of other firms in the network without participating in the grand coalition and sharing security costs. Such free-rider issues have also been identified in other contexts to hinder cooperation and stability of the grand coalition in other partition function form games (see, e.g., Yi, 1997).

Proposition 7. *The grand coalition in the interdependent security cost-sharing game under public information,* $(N,\widehat{c})$*, is not, in general, stable to defections.*

We now, however, show that the agreeable allocation can be extended to the public information setting while retaining several of its desirable properties. Notably, we prove that, analogous to Theorem 6, the public information version of the agreeable allocation, when it exists, satisfies individual rationality, a weaker notion of stability wherein each player is better off in the grand coalition (i.e., with full cooperation) as compared to the independent coalitions (i.e., no-cooperation) scenario.

6.1 | Agreeable allocation with public information

Again, for ease of exposition, we restrict our attention to networks where all firms are secured in the grand coalition. We recursively define a finite family of mutually exclusive sets $\mathcal{T} = \{\mathcal{T}_1, \ldots, \mathcal{T}_\ell\}$ of players in the network where $\mathcal{T}_1 = \{i \in N : \widehat{\Upsilon}^i_{\{i\};\rho_1} = 1\}$, where ρ_1 corresponds to the independent coalition structure. For $k \geq 1$, we then define $\mathcal{T}_{2k}$ and $\mathcal{T}_{2k+1}$ recursively as follows, where $\overline{\mathcal{T}_k} = \mathcal{T}_1 \cup \cdots \cup \mathcal{T}_k$. Further, the coalition structure ρ_{k+1} contains the coalition $\overline{\mathcal{T}_k}$ and all other players in $N \setminus \overline{\mathcal{T}_k}$ are in independent coalitions. Also, recall that $\widehat{\Upsilon}^i_{S;\rho}$ is the equilibrium security state of player $i \in S$ with the coalition structure ρ in the public information model whereas Υ^i_S is the coalition-optimal security state of $i \in S$ in the private information setting.

$$\mathcal{T}_{2k} = \left\{ i \in N \setminus \overline{\mathcal{T}_{2k-1}} : \Upsilon^i_{\overline{\mathcal{T}_{2k-1}} \cup \{i\}} = 1 \right\}, \tag{10}$$

$$\mathcal{T}_{2k+1} = \left\{ i \in N \setminus \overline{\mathcal{T}_{2k}} : \widehat{\Upsilon}^i_{\overline{\mathcal{T}_{2k}} \cup \{i\};\rho_{2k+1}} = 1 \right\}. \tag{11}$$

$\mathcal{T}_1$ contains players that are secured under the independent coalition structure. That is, in the equilibrium outcome obtained from Algorithm 3, these players are secured. $\mathcal{T}_2$ contains players who, if they are secured, save the costs of extrinsic security for players in $\mathcal{T}_1$ and bestow a direct positive externality to the players in $\mathcal{T}_1$ that outweighs their own cost of security. Thus, for the players in $\mathcal{T}_1 \cup \mathcal{T}_2$, it is optimal in the private information model as well to secure themselves. Further, there will be players in $\mathcal{T}_3$ for whom it is individually rational to secure themselves conditional upon players in $\mathcal{T}_1$ and $\mathcal{T}_2$ being in a coalition together, $\mathcal{T}_1 \cup \mathcal{T}_2$. Successive sets of players are identified iteratively. Note that these families of sets are constructed in a very similar manner as in the private information model. The only distinction arises in (11) from observing that in a public information model, the formation of each new coalition may also trigger a change in the security actions of other players who can respond to this.

Suppose there exists $\ell \in \mathbb{Z}$ such that $\overline{\mathcal{T}_\ell} = N$, then the recursive procedure generating the family of sets terminates. Again, it is possible in certain networks and associated cost parameter vectors for no $\ell \in \mathbb{Z}$ to exist such that $\overline{\mathcal{T}_\ell} = N$. In these cases, consequently, no agreeable allocation will exist. Unlike in the private information setting where a closed form expression for the agreeable allocation is derived, the agreeable allocation under public information $\hat{x}$ is obtained by Algorithm 5 provided in Section EC.3, which takes in the family of sets $\mathcal{T}$ as an input.

Theorem 9. *The agreeable allocation under public information, $\hat{x}$, when it exists, is (i) individually rational, (ii) polynomial-time computable, and (iii) bilaterally implementable. Further, it also satisfies (iv) symmetry and the (v) null player property.*

Therefore, while the agreeable allocation cannot guarantee that the grand coalition is stable to defections by subsets of players (indeed no cost-sharing allocation can), it still satisfies a weaker notion of stability. It ensures that all players will prefer to remain in the grand coalition structure ρ^* rather than in the independent coalition structure. Further, we interestingly find that the public information version of the agreeable allocation exists if and only if the agreeable allocation as defined in the private information setting exists.

Corollary 2. *For a given network $\mathbb{G} = (N, A)$ and associated security cost parameters, the agreeable allocation under public information $\hat{x}$ exists if and only if the agreeable allocation under private information x^* exists.*

Here, we briefly comment on some main implications of our analysis of the general partial information model in Section EC.4. First, we demonstrate that the agreeable allocation can be naturally extended to the partial information model thereby generalizing Theorem 9. Therein, we observe that Corollary 2 also generalizes and the existence of the agreeable allocation is not contingent on the informational assumption in the network. Finally, and importantly, we clarify that even in the presence of partial public information in the network, the grand coalition may be unstable and that if the grand coalition is unstable with a certain level of public information in the network, it remains unstable at higher levels of information provisioning in the network.

7 | QUASI-HOMOGENEOUS NETWORKS

The chief deficiency of the agreeable allocation, under all informational assumptions is that, in general, depending on the structure of the interfirm network, or the associated security costs, it may not exist. To the extent that an agreeable allocation is viewed as desirable for its fairness, bilateral implementability, and other properties as documented in Theorem 6 and Theorem 9, this offers a rationale for when interfirm networks will find it challenging to cooperatively secure themselves. In order to examine the role of the network structure on the existence of the agreeable allocation, we now consider *quasi-homogeneous* networks $\mathbb{G}$ as networks wherein the costs of securing against intrinsic risks for firm i, θ_i, are identical for all firms. Similarly, we also assume costs of securing against extrinsic risks, ξ_{ij}, are identical across all links in the network, and the expected penalties faced by players in the event of a realized risk are also equal. Formally, a network $\mathbb{G}$ is said to be quasi-homogeneous if $\theta_i = \theta$ and $L_i = L$ for all $i \in V$ and $\xi_{ij} = \xi$ for all $(i,j) \in A$.

Analyzing quasi-homogeneous networks permits us to isolate the effects of the network structure on the existence of the agreeable allocation. A priori, it is qualitatively unclear what

the role of network structure would be on the existence of the bilaterally implementable agreeable allocation. For instance, denser networks can render it easier for efficient and stable cost-sharing arrangements to be bilaterally implementable since there are more bilateral links. However, denser networks may also result in wider positive externalities to securing oneself necessitating multilateral cooperation.

We now introduce some graph-theoretic definitions that aid us in identifying when quasi-homogeneous networks admit and do not admit an agreeable allocation of security costs. We define a *k-core* of network $\mathbb{G}$ as an induced subgraph $\mathbb{H}$ of $\mathbb{G}$ such that the in-degree of all nodes in $\mathbb{H}$ is at least k.[9] Then, a (k,ℓ)-*core* is a k-core $\mathbb{H}$ of $\mathbb{G}$ such that, if ℓ denotes the maximum out-degree of a node in $\mathbb{H}$ to the nodes in $\mathbb{G}\backslash\mathbb{H}$, then $k > \ell$. Therefore, while a k-core is a sufficiently dense induced subgraph, a (k,ℓ)-*core* is an induced subgraph that is sufficiently dense internally and simultaneously sparse in its connections with other nodes in the graph.

Theorem 10. *Consider a quasi-homogeneous network* $\mathbb{G}$ *with security cost parameters given by* L, θ, *and* ξ.

(*i*) $\mathbb{G}$ *admits an agreeable allocation if* $\mathbb{G}$ *does not contain a* k*-core where* $k = \left\lceil \frac{L-\theta}{\xi} \right\rceil$.

(*ii*) $\mathbb{G}$ *does not admit an agreeable allocation if* $\mathbb{G}$ *contains a* (k,l)-core *where* $k = \ell + \left\lceil \frac{L-\theta}{\xi} \right\rceil$.

The two parts of Theorem 10 provide distinct sufficient and necessary conditions, respectively, for the existence of the agreeable allocation in quasi-homogeneous networks. From a descriptive standpoint, it implies qualitatively that the agreeable allocation is guaranteed to exist in (quasi-homogeneous) networks so long as they are not sufficiently locally dense. This refines our earlier intuition on the role of interfirm network structure on the existence of the agreeable allocation. Further, in graphs that contain sufficiently dense and sufficiently local clusters, the agreeable allocation is guaranteed to not exist.

8 | NUMERICAL CASE STUDY

We now present a case study analyzing the feasibility of cost-sharing mechanisms to sustain network-wide cooperative security in real-world interfirm networks that can face interdependent risks. Specifically, we use the Refinitiv SDC Alliance database to extract all alliances in the food manufacturing sector formed between 2006 to 2020. The database contains 2339 alliances formed between 3073 unique firms in our industry of interest. Typically, these are bilateral alliances formed between two firms, while, on occasion, alliances are formed between three or more firms. For example, one of the alliances in the database is between Optibiotix Health Plc, a biotechnology company that manufactures SlimBiome, a weight management supplement, and John Morley (Importers) Ltd, which manufactures prepared perishable foods. Optibiotix Health would supply the weight management supplement to be included in prepared muesli packs manufactured by John Morley Ltd within the United Kingdom. In this example, the presence of an interdependent risk is evident. Over time, larger networks of alliances arise and we identify 792 distinct interfirm networks. Of these, the largest connected network of firms contains 1092 nodes. The other networks are smaller, and we remove all networks consisting of only two firms since these networks trivially permit bilaterally implementable cost-sharing mechanisms. We in fact restrict our attention to alliance networks that are of size at least five, and we obtain exactly 50 such alliance networks.[10] We depict two of these networks in Figure 2.

We leverage the algorithmic results obtained in previous sections to numerically test whether the agreeable allocation exists, and when it exists, compute the network-wide security cost apportioned by the allocation. These results are meant to be illustrative since the existence of the agreeable allocation naturally depends on the precise security cost parameter specifications. However, the security cost parameters and the penalties are simulated in a systematic manner. Across all simulated networks, we set the parameter $\theta_i \sim U[15,25]$ for all firms i, and for all links between firms i and j, (i,j), $\xi_{ij} \sim U[3,5]$. Further, for all i, $L_i \sim U[17+\delta_i, 23+\delta_i]$, where $\delta_i = |N^-|$. That is, we assume that firms with more partners are larger firms and thus also likely to incur higher reputation costs. Based on 1000 simulated runs for each of the 50 alliance networks, we make the following observations.

First, we observe that in 56.7% of the simulated networks, the agreeable allocation exists. In contrast, in only 0.79% of the simulated networks, the Shapley value based security cost-sharing allocation is of the form given by Theorem 4 and, hence, bilaterally implementable. This, in conjunction with the straightforward implementation mechanism described in Section 5, demonstrates the practical relevance of our proposed security cost-sharing allocation. Second, we find, interestingly, that the alliance network permitting the agreeable allocation to exist with the highest likelihood of 74.3%, is a star network. Finally, we observe that the networks which rarely permit the existence of the agreeable allocation, in only 2.6% and 4% of the simulations, respectively, are both completely connected networks, that is, cliques of size six. This lends further evidence in support of Theorem 10 that densely connected networks preclude the existence of the agreeable allocation.

In the above numerical experiment, the cost parameters for all nodes in a network were drawn from the same distributions. However, in real-world networks, there is usually a significant asymmetry in the penalties incurred by firms in case of a realized risk. Consumer-facing firms typically incur substantially larger penalties than others. To incorporate this in our simulation, we obtain the Standard Industrial Classification (SIC) codes of the firms from the SDC database. We then denote firms in the retail industry (with a SIC code in the range of 5200–5999) as consumer-facing firms. Of the 3073 unique firms in our dataset, we identify 154 such (potentially) consumer-facing firms. In our second

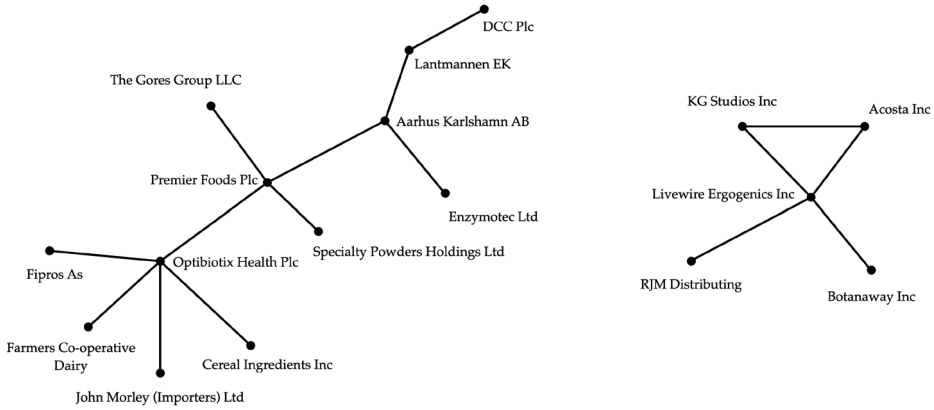

FIGURE 2 Examples of alliance networks in the food manufacturing sector.

numerical experiment, we simulate the cost parameter L_i for a firm i such that a consumer-facing firm faces a larger expected penalty and the expected penalty decays exponentially with the distance from the consumer, that is, $L_i = L_0/c_0^{d_i}$, where $d_i = 0$ if i is a consumer-facing firm, L_0 is the expected penalty it faces, and c_0 is a constant.[11] We again perform 1000 simulation runs for each of the 50 alliance networks. Each network is then compared against a benchmark simulation wherein the penalties of all firms are drawn from the same uniform distribution with an expected penalty given by $(\sum_{i\in N} L_0/c_0^{d_i})/N$. This allows us to comment on the role of cost asymmetry on the existence of the agreeable allocation vis-á-vis the bilateral implementability of the Shapley mechanism. For our chosen parameter values, we find that in the benchmark network simulations, the Shapley value nearly always coincides with the agreeable allocation and is bilaterally implementable for all of the 50 networks. However, with asymmetric penalties, the Shapley value is bilaterally implementable only in 34.47% of the simulations. For 15 of the 50 networks, it was never bilaterally implementable across all 1000 runs. In contrast, the bilaterally implementable agreeable allocation exists in 71.35% of the simulated networks. Across various choices of L_0 and c_0, we recover qualitatively identical results. In summary, in real-world networks with cost asymmetries, despite the nonexistence of the agreeable allocation in certain instances, the practical advantage[12] of the agreeable allocation in terms of its bilaterally implementability over the Shapley mechanism is further underscored.

9 | CONCLUDING REMARKS

Networked firms are exposed to a variety of interdependent, or contagion, risks such as supply chain contamination, deliberate adulteration, or cybersecurity threats and data breaches. The fundamental distinction that sets apart these risks from other types of risks faced by firms is their transferable nature. In this paper, we develop a network model to study the cooperative management of interdependent risks by networked firms.

The network-wide cooperative security strategy in our interdependent risk model can be computed in polynomial time via a minimum-weight cut network flow algorithm. Assuming that the security costs and actions are private information known only to the respective players, we find that firms have a clear incentive to cooperate and that there exist stable security cost-sharing mechanisms that can sustain network-wide cooperation. However, in the presence of public information, we find that, in general, there do not exist cost-sharing mechanisms that can ensure the stability of the grand coalition. Thus, it appears that interdependence of network security is alone insufficient to sustain network-wide cooperation.

Introducing the notion of bilateral implementability, we uncover a fundamental trilemma between stability, fairness, and implementability of network security cost-sharing mechanisms. We then develop a novel cost-sharing mechanism, the *agreeable allocation*, which attempts to balance the three notions. The agreeable allocation, when it exists, satisfies notions of stability, is formalizably fair, easily computable, and is also implementable via a series of bilateral cost-sharing agreements. However, the agreeable allocation may not always exist. This, we argue, once again, demonstrates that, although cost-sharing mechanisms belonging to the core can be identified, sustaining network-wide security cooperation can still be challenging and, therefore, may not always be possible in practice. We then construct δ-agreeable allocations that satisfy the general notion of $(\delta + 1)$-implementability which permits firms that are not direct partners to also enter into cost-sharing agreements if they are at a distance of at most δ from each other in the network. As δ increases, the δ-agreeable allocation is more likely to exist. However, as δ increases, we also expect the coordination challenges to increase thereby highlighting a fundamental trade-off.

Moreover, to study the role of network structure on the existence of the agreeable allocation, we consider quasi-homogeneous networks (i.e., networks with homogeneous costs of security and expected penalties in case of realized risk) and find that networks without sufficiently dense clusters admit an agreeable allocation. Whereas, networks

containing sufficiently dense and local clusters do not permit an agreeable allocation of network-wide security costs. Finally, using the SDC alliance database, we extract all alliances formed in the food manufacturing sector between 2006 and 2020. With numerical experiments and simulated cost parameters, we argue the practical feasibility and relevance of employing the agreeable allocation as a bilateral security cost-sharing mechanism in real-world alliances to sustain network-wide cooperative security against interdependent risks.

This work develops, to the best of our knowledge, for the first time, an economic theory of cooperative security against interdependent risks in networks. However, we acknowledge several limitations and open problems arising from our study.

Limitations

First, for instance, the question of the general existence (or nonexistence) of a bilaterally implementable and stable cost-sharing mechanism remains open. Second, and crucially, in this paper, we consider interfirm networks characterized by repeated and ongoing interactions between firms. Thus, a vulnerable firm is nearly certain to transfer risks to its partner firms if the partner firms do not secure the corresponding link. A richer model of interdependent security would allow for a stochastic transmission and propagation of risk in the network. However, this richer stochastic model of interdependent network security is challenging to analyze. Particularly, the characterization of cooperative security strategies in this stochastic model of interdependent security is a nontrivial problem. Finally, we assume that the considered networks are static whereas, in reality, networks tend to change dynamically, with new alliances being formed, and existing alliances being broken over time. Bilaterally implementable cost-sharing mechanisms, in particular, may be well-suited to sustain cooperation in dynamic alliances, as we have noted earlier.

ACKNOWLEDGMENTS

We thank Amitabh Basu of Johns Hopkins University for helpful private communication pertaining to Theorem 10. We also thank the Department Editor, the Associate Editor, and two anonymous referees for their constructive suggestions that led to a much improved paper.

ORCID

Sriram Sankaranarayanan https://orcid.org/0000-0002-1224-5422

ENDNOTES

[1] Relatedly, Dawande and Qi (2021), in a review of recent research on socially responsible operations management, note that "a topic that has not received much attention yet is the design of cooperative strategies among stakeholders in different tiers of a supply chain to collectively ensure socially responsible actions across the supply chain … the utilities of different players from actions such as auditing, inspections, and testing become interconnected in a complex manner. Consequently, the sharing of costs in a fair manner to incentivize cooperation across tiers becomes challenging."

[2] The terms *agents*, *firms*, and *players* are used interchangeably in this paper.

[3] In the interdependent security literature, intrinsic and extrinsic risks are sometimes referred to as *direct* and *indirect* risks, respectively.

[4] Relatedly, Roughgarden (2010) observes "(A) complexity-theoretic hardness result can diminish the predictive interpretation of an equilibrium concept and suggests more tractable alternatives […] In a practical design context, it is obvious that a mechanism that is actually implemented had better be computationally tractable to run, like the deferred acceptance algorithm, and also easy to play, in the sense that participants should not need to perform difficult computations."

[5] Furthermore, a purely cooperative-game theoretic approach to cost-sharing problems on occasion faces some criticism, as, for example, in Feng et al. (2021), of providing "no implication for implementation in terms of how firms interact in the network and how financial payments are made among the firms."

[6] P is said to be minimal if it is a coalitionally rational security set for i but no subset of P is.

[7] In the general partial information model analyzed in Section EC.4, firm i only adopts a worst-case belief for firms whose information is private, that is, for $j \in N \setminus \mathcal{P}$, $\sigma_{ji} = 0$ whereas for $j \in \mathcal{P}$, i forms an accurate belief, $\sigma_{ji} = \sigma_j$.

[8] Our equilibrium selection procedure bears resemblance and is motivated by the level-k approach (Stahl & Wilson, 1995) which yields sufficient conditions for an equilibrium.

[9] Conventionally, k-cores are defined on undirected graphs. Herein, we consider a natural analogue for directed graphs.

[10] The largest alliance network comprises 1092 firms and 2624 partnerships (i.e., arcs). The other 49 alliance networks are smaller and qualitatively bear structural similarities containing an average of 6.79 nodes (a median of six nodes) and 14.28 arcs (a median of 12). The average degree of each node across the 50 alliance networks (i.e., the average number of partners for a firm) is 2.13. We also observe that 28 of these 50 alliance networks are *trees*.

[11] L_0 was chosen to be 409.6 ($=2^{11}/5$) and $c_0 = 2$ for the results reported here.

[12] Not surprisingly, we also observed a substantial advantage in terms of the computational time required to obtain the agreeable allocation in comparison to the Shapley value.

REFERENCES

Acemoglu, D., Malekian, A., & Ozdaglar, A. (2016). Network security and contagion. *Journal of Economic Theory*, *166*, 536–585.

ASEM. (2013). The Vienna Declaration. https://aseminfoboard.org/wp-content/uploads/2022/10/10th_ASEM_DGs_-_FINAL_VIENNA_DECLARATION_yQ7Aiyl.pdf

Aydinliyim, T., & Vairaktarakis, G. (2010). Coordination of outsourced operations to minimize weighted flow time and capacity booking costs. *Manufacturing & Service Operations Management*, *12*(2), 236–255.

Babich, V., & Tang, C. S. (2012). Managing opportunistic supplier product adulteration: Deferred payments, inspection, and combined mechanisms. *Manufacturing & Service Operations Management*, *14*(2), 301–314.

Blaettchen, P., Calmon, A. P., & Hall, G. (2021). *Traceability technology adoption in supply chain networks*. arXiv. https://doi.org/10.48550/arXiv.2104.14818

Caro, F., Chintapalli, P., Rajaram, K., & Tang, C. S. (2018). Improving supplier compliance through joint and shared audits with collective penalty. *Manufacturing & Service Operations Management*, *20*(2), 363–380.

Caro, F., Lane, L., & de Tejada Cuenca, A. S. (2021). Can brands claim ignorance? Unauthorized subcontracting in apparel supply chains. *Management Science*, *67*(4), 2010–2028.

Chan, H., Ceyko, M., & Ortiz, L. E. (2012). *Interdependent defense games: Modeling interdependent security under deliberate attacks*. arXiv. https://doi.org/10.48550/arXiv.1210.4838

Chan, H., & Ortiz, L. E. (2014). Computing Nash equilibria in generalized interdependent security games. *Advances in Neural Information Processing Systems*, *27*, 2735–2743.

Chen, J., Qi, A., & Dawande, M. (2020). Supplier centrality and auditing priority in socially responsible supply chains. *Manufacturing & Service Operations Management*, *22*(6), 1199–1214.

Chen, R. R., & Yin, S. (2010). The equivalence of uniform and Shapley value-based cost allocations in a specific game. *Operations Research Letters*, *38*(6), 539–544.

Dawande, M., & Qi, A. (2021). Auditing, inspections, and testing for social responsibility in supply networks. In *Responsible Business Operations* (pp. 243–259). Springer.

Dhingra, V., & Krishnan, H. (2021). Managing reputation risk in supply chains: The role of risk sharing under limited liability. *Management Science*, *67*(8), 4845–4862.

Dong, L., Qiu, Y., & Xu, F. (2022). Blockchain-enabled deep-tier supply chain finance. *Manufacturing & Service Operations Management*. Advance online publication. https://doi.org/10.1287/msom.2022.1123

Ergun, Ö., Gui, L., Heier Stamm, J. L., Keskinocak, P., & Swann, J. (2014). Improving humanitarian operations through technology-enabled collaboration. *Production and Operations Management*, *23*(6), 1002–1014.

Fang, X., & Cho, S.-H. (2020). Cooperative approaches to managing social responsibility in a market with externalities. *Manufacturing & Service Operations Management*, *22*(6), 1215–1233.

Feng, Q., Li, C., Lu, M., & Shanthikumar, J. G. (2021). Implementing environmental and social responsibility programs in supply networks through multiunit bilateral negotiation. *Management Science*, *68*(4), 2579–2599.

Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., & Yariv, L. (2010). Network games. *The Review of Economic Studies*, *77*(1), 218–244.

Goldberg, A. V., & Tarjan, R. E. (1988). A new approach to the maximum-flow problem. *Journal of the ACM*, *35*(4), 921–940.

Gopalakrishnan, S., Granot, D., & Granot, F. (2021a). Consistent allocation of emission responsibility in fossil fuel supply chains. *Management Science*, *67*(12), 7637–7668.

Gopalakrishnan, S., Granot, D., Granot, F., Sošić, G., & Cui, H. (2021b). Incentives and emission responsibility allocation in supply chains. *Management Science*, *67*(7), 4172–4190.

Granot, D., & Sošić, G. (2003). A three-stage model for a decentralized distribution system of retailers. *Operations Research*, *51*(5), 771–784.

Gui, L., Atasu, A., Ergun, Ö., & Toktay, L. B. (2018). Design incentives under collective extended producer responsibility: A network perspective. *Management Science*, *64*(11), 5083–5104.

Hafalir, I. E. (2007). Efficiency in coalition games with externalities. *Games and Economic Behavior*, *61*(2), 242–258.

Heal, G., & Kunreuther, H. (2007). Modeling interdependent risks. *Risk Analysis: An International Journal*, *27*(3), 621–634.

Huang, L., Song, J.-S. J., & Swinney, R. (2020). *Managing social responsibility in multitier supply chains*. SSRN. https://ssrn.com/abstract=2837332

Huang, X., Boyacı, T., Gümüş, M., Ray, S., & Zhang, D. (2016). United we stand or divided we stand? Strategic supplier alliances under order default risk. *Management Science*, *62*(5), 1297–1315.

Kearns, M. J., & Ortiz, L. E. (2003). Algorithms for interdependent security games. *Advances in Neural Information Processing Systems*, *16*, 561–568.

Kemahlıoğlu-Ziya, E., & Bartholdi III, J. J. (2011). Centralizing inventory in supply chains by using Shapley value to allocate the profits. *Manufacturing & Service Operations Management*, *13*(2), 146–162.

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

Laszka, A., Felegyhazi, M., & Buttyan, L. (2014). A survey of interdependent information security games. *ACM Computing Surveys (CSUR)*, *47*(2), 1–38.

Lawrence, F. (2013). Horsemeat scandal: The essential guide. The Guardian, February 15.

Lee, H.-H., & Li, C. (2018). Supplier quality management: Investment, inspection, and incentives. *Production and Operations Management*, *27*(2), 304–322.

Leng, M., & Parlar, M. (2009). Allocation of cost savings in a three-level supply chain with demand information sharing: A cooperative-game approach. *Operations Research*, *57*(1), 200–213.

Levi, R., Singhvi, S., & Zheng, Y. (2020). Economically motivated adulteration in farming supply chains. *Management Science*, *66*(1), 209–226.

McAfee. (2015). Cloud security: Target breach reveals risk of business partners. *April 28*. https://www.mcafee.com/blogs/enterprise/cloud-security/target-breach-reveals-risk-of-business-partners/

Mu, L., Dawande, M., Geng, X., & Mookerjee, V. (2016). Milking the quality test: Improving the milk supply chain under competing collection intermediaries. *Management Science*, *62*(5), 1259–1277.

Mu, L., Dawande, M., & Mookerjee, V. (2019). Shaping the values of a milk cooperative: Theoretical and practical considerations. *Production and Operations Management*, *28*(9), 2259–2278.

Nagarajan, M., & Sošić, G. (2008). Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *European Journal of Operational Research*, *187*(3), 719–745.

Plambeck, E. L., & Taylor, T. A. (2016). Supplier evasion of a buyer's audit: Implications for motivating supplier social and environmental responsibility. *Manufacturing & Service Operations Management*, *18*(2), 184–197.

Pollach, I. (2011). Online privacy as a corporate social responsibility: An empirical study. *Business Ethics: A European Review*, *20*(1), 88–102.

Rodríguez-Pereira, J., Balcik, B., Rancourt, M.-È., & Laporte, G. (2021). A cost-sharing mechanism for multi-country partnerships in disaster preparedness. *Production and Operations Management*, *30*, 4541–4565.

Roughgarden, T. (2010). Computing equilibria: A computational complexity perspective. *Economic Theory*, *42*(1), 193–236.

Shapley, L. S. (1953). A value for n-person games. *Annals of Mathematics Study*, *28*, 307–317.

Shapley, L. S. (1971). Cores of convex games. *International Journal of Game Theory*, *1*(1), 11–26.

Snyder, L. V., Scaparra, M. P., Daskin, M. S., & Church, R. L. (2006). Planning for disruptions in supply chain networks. In *Models, methods, and applications for innovative decision making* (pp. 234–257). Informs.

Sošić, G. (2006). Transshipment of inventories among retailers: Myopic vs. farsighted stability. *Management Science*, *52*(10), 1493–1508.

Stahl, D. O., & Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, *10*(1), 218–254.

Tian, F., Sošić, G., & Debo, L. (2020). Stable recycling networks under the extended producer responsibility. *European Journal of Operational Research*, *287*(3), 989–1002.

Westerink-Duijzer, L. E., Schlicher, L. P., & Musegaas, M. (2020). Core allocations for cooperation problems in vaccination. *Production and Operations Management*, *29*(7), 1720–1737.

Yi, S.-S. (1997). Stable coalition structures with externalities. *Games and economic behavior*, *20*(2), 201–237.

Zhang, H., Aydin, G., & Parker, R. P. (2022). Social responsibility auditing in supply chain networks. *Management Science*, *68*(2), 1058–1077.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Gopalakrishnan, S., & Sankaranarayanan, S. (2023). Cooperative security against interdependent risks. *Production and Operations Management*, 1–17. https://doi.org/10.1111/poms.14047