

Dziubiński, Marcin; Goyal, Sanjeev

Article

How do you defend a network?

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Dziubiński, Marcin; Goyal, Sanjeev (2017) : How do you defend a network?, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 12, Iss. 1, pp. 331-376,
<https://doi.org/10.3982/TE2088>

This Version is available at:

<https://hdl.handle.net/10419/197191>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/3.0/>

How do you defend a network?

MARCIN DZIUBIŃSKI

Institute of Informatics, Faculty of Mathematics, Informatics, and Mechanics, Warsaw University

SANJEEV GOYAL

Faculty of Economics and Christ's College, University of Cambridge

Modern economies rely heavily on their infrastructure networks. These networks face threats ranging from natural disasters to human attacks. As networks are pervasive, the investments needed to protect them are very large; this motivates the study of targeted defense. What are the “key” nodes to defend to maximize functionality of the network? What are the incentives of individual nodes to protect themselves in a networked environment and how do these incentives correspond to collective welfare?

We first provide a characterization of optimal attack and defense in terms of two classical concepts in graph theory: *separators* and *transversals*. This characterization permits a systematic study of the *intensity of conflict* (the resources spent on attack and defense) and helps us identify a new class of networks—*windmill graphs*—that minimize conflict.

We then study security choices by individual nodes. Our analysis identifies the externalities and shows that the welfare costs of decentralized defense in networks can be very large.

KEYWORDS. Infrastructure, costs of conflict, windmill graph, attack, defense.

JEL CLASSIFICATION. C7, D1, D8.

1. INTRODUCTION

Our nation's critical infrastructure is crucial to the functioning of the American economy... (It) is increasingly connected and interdependent and protecting it and enhancing its resilience is an economic and national security imperative (Department of Homeland Security 2012).

Marcin Dziubiński: m.dziubinski@mimuw.edu.pl

Sanjeev Goyal: sg472@cam.ac.uk

We thank the co-editor and three anonymous referees for comments that have greatly helped improve the exposition of the paper. Thanks are also due to Diego Cerdeiro, Syngjoo Choi, Amrita Dhillon, Matt Elliott, Julien Gagnon, Timo Hiller, Prem Chand Jain, Peter Kondor, Michael König, Botond Kőszegi, Francesco Nava, Anja Prummer, Pauline Rutsaert, Adam Szeidl, and numerous seminar participants for comments. Both authors thank the European Research Area Complexity Net for financial support. Marcin Dziubiński was supported by the Strategic Resilience of Networks project realized within the Homing Plus program of the Foundation for Polish Science and was co-financed by the European Union from the Regional Development Fund within Operational Programme Innovative Economy (grants for innovation). Sanjeev Goyal acknowledges financial support from a Keynes Fellowship and the Cambridge INET Institute.

Infrastructure networks—highways, aviation, shipping, pipelines, train systems, and posts—are a vital part of the modern economy. These networks face a variety of threats ranging from natural disasters to human attacks. The latter may take a violent form (guerrilla attacks, attacks by an enemy country, and terrorism) or a nonviolent form (as in political protest that blocks transport services).¹ A network can be made robust to such threats through additional investments in equipment and in personnel. As networks are pervasive, the investments needed could be very large; this motivates the study of targeted defense. What are the “key” parts of the network that should be protected to ensure maximal functionality? As defense is often a choice made by individual actors, we also wish to understand the relation between network structure and decentralized incentives. This paper develops a model to study these questions.

Consider a given infrastructure network consisting of nodes and links. The defender chooses to protect “nodes” of the network against damage/attacks; protecting a node is costly. Protection includes investments in security personnel, in training, in equipment, and in cybersecurity. These protection measures typically take time to implement and so we focus on *ex ante* investments in protection. We suppose that a defended node is immune to attack whereas an undefended node is eliminated by attack (along with all its links). The initial network, the defense, and the attack together yield a set of surviving nodes and links—the residual network. The defender chooses a defense strategy that maximizes the value of the residual network, net of the costs of defense.

Our model covers two scenarios. The first is that of an intelligent adversary who seeks to damage components and disrupt the flows in the network. The second is that of a natural threat: facing such a threat, the defender focuses on the worst case scenario. In both cases the defender looks for the “maximin” solution. For expositional simplicity, we use the language of an intelligent adversary throughout. We study a game between a defender and an adversary and analyze the subgame perfect equilibrium of this game.

We consider network payoff functions in which the value to the defender of a network is component additive, and the payoff from each component is increasing and convex in the size of the component.² The convexity of value in component size is key to the appeal of connectivity in networks.

We begin with a study of optimal defense. **Proposition 2** characterizes optimal defense and attack. Optimal attack targets two types of nodes: those that fragment the network into distinct components (the *separators*) and those that simply reduce the size of components (the *reducing attacks*). As payoffs are convex in component size, separators are particularly attractive targets for attack (as their elimination disconnects components). Anticipating this attack, optimal defense targets nodes that block the separators and reduce attacks. A set of nodes that block a collection of separators is referred to

¹For an introduction to network based conflict, see Arquilla and Ronfeldt (2001) and Zhu and Levinson (2011); for news coverage of the effects of natural disasters and human attacks on infrastructure networks, see Eun (2010), Kliesen (1994), and Pandey (2011) and Luft (2005).

²This specification is consistent with Metcalfe’s law (network value is proportional to the square of the number of nodes) and Reed’s law (network value is exponentially increasing in the number of nodes). It is also in line with the large theoretical literature on network externalities (Katz and Shapiro 1985, Farrell and Saloner 1986) and network economics (Bala and Goyal 2000, Jackson and Wolinsky 1996). One way to define network value is the number of pairs of nodes connected (directly or indirectly) in the network. This is a special case of our value function.

as a *transversal*. We prove that optimal defense either targets a minimal transversal or protects all nodes. Figures 3 and 4 illustrate these concepts.³

This characterization result allows us to study the relation between networks and conflict more closely. We find that the size of defense and attack are both nonmonotonic in the cost of attack; even more surprisingly, the size of defense and the payoff of the defender may fall with the addition of links in the network (Proposition 3).

We then turn to the *intensity of conflict*: this is the sum of expenditures of defense and attack. For a given configuration of costs of defense and attack, we derive the minimal intensity of conflict and then describe the networks that sustain it (Proposition 4). We then demonstrate that network architecture can create very large variations in the intensity of conflict. A feature of minimal conflict is that there is a single active player. We next discuss circumstances under which both players devote resources to conflict in equilibrium.

An important insight of the analysis is the optimality of *strategic exposure*: the defender may find it optimal to leave unprotected a key node (the elimination of which disconnects the network) and instead to protect an alternative, larger, set of nodes. We refer to this as the *queen sacrifice*. This leads us to identify a class of networks—*windmill graphs*—that minimize conflict and are also attractive for the defender. Figure 7 presents these networks.

In many situations, security decisions are made at the local level, e.g., individual airports choose their own security checks. This motivates the study of decentralized security.⁴ Individual nodes care about surviving an attack and about being part of a large connected network. Observe that to block a separator it is sufficient for one node in the separator to protect itself. So, in the game among the nodes, defense choices within a separator are strategic substitutes. But for the network to remain connected, all separators must be blocked. Therefore a node will protect itself only if other separators are being blocked: thus defense choices also exhibit strategic complementarity. Proposition 5 shows that decentralized security choices can be characterized in terms of separators and transversals of the network. Finally, we demonstrate that a combination of incentive and coordination issues may lead to very large costs of decentralization.

Our paper contributes to the economic study of networks. The research on networks has been concerned with the formation, structure, and functioning of social and economic networks (Goyal 2007, Jackson 2008, and Vega-Redondo 2007). The problem of key players has traditionally been studied in terms of Bonacich centrality, betweenness, eigenvectors, and degree centrality; see, e.g., Bala and Goyal (2000), Ballester et al. (2006), Choi et al. (forthcoming), DeMarzo et al. (2003), Elliott and Golub (2013), Galeotti et al. (2010), and Golub and Jackson (2010). Our paper suggests that for the problem of attack and defense, the key players are nodes that lie in separators and transversals. These nodes are typically distinct from nodes that maximize familiar notions of centrality. Appendix B discusses this distinction in detail. Thus the principal

³Appendix C provides a detailed application of the concepts to well known families of networks (trees, core-periphery, interlinked stars).

⁴For an early contribution on interdependent security, see Kunreuther and Heal (2003).

contribution of our paper is to introduce two classical concepts from graph theory into economics and show how they address a problem of practical importance.

Individual defense is a public good, and so this conceptual contribution is also relevant for the study of games on networks more generally. [Bramoullé and Kranton \(2007\)](#) draw attention to maximal independent sets. By contrast, our work brings out the role of minimal transversal of the separators. These sets are generally different from maximal independent sets.⁵

Our paper also contributes to the literature on network defense; see, e.g., [Bier et al. \(2007\)](#), [Baccara and Bar-Isaac \(2008\)](#), [Acemoglu et al. \(2013\)](#), [Dziubiński and Goyal \(2013\)](#), [Goyal and Vigier \(2014\)](#), [Clark and Konrad \(2007\)](#), and [Kovenock and Roberson \(2012\)](#). To the best of our knowledge, our results on the role of separators and transversals in network conflict are novel, relative to the existing body of work. In particular, we note that the earlier work by [Dziubiński and Goyal \(2013\)](#) and [Goyal and Vigier \(2014\)](#) focuses on optimal design and defense. In these papers, the optimal network takes on a very simple form—it is a star—and so the optimal defense takes on a correspondingly simple structure: protect the central hub node. By contrast, in the present paper the network is exogenous and arbitrary: this is a much broader problem and requires new conceptual tools.

We note that the problem of network defense has traditionally been studied in operations research, electrical engineering, and computer science; see, e.g., [Alpcan and Başar \(2011\)](#), [Aspnes et al. \(2006\)](#), [Smith \(2008\)](#), and [Grötschel et al. \(1995\)](#). In an early paper, [Cunningham \(1985\)](#) looks at the problem of network design and defense with conflict on links. Relative to this literature, the novelty of our paper lies in the study of intensity of conflict and the externalities that arise in decentralized defense.

The rest of the paper is organized as follows. [Section 2](#) presents the model of defense and attack. [Section 3](#) introduces the main concepts and provides a characterization of equilibrium defense and attack. It also contains the study of comparative statics, active conflict, and conflict intensity. [Section 4](#) takes up the case of decentralized defense. [Section 5](#) concludes. All proofs are presented in [Appendix A](#). [Appendix B](#) analyzes the relation between key nodes to attack and defend and other notions of centrality. [Appendix C](#) illustrates the notions of separators and transversals in well known families of networks such as core–periphery networks, trees, interlinked stars, and bipartite graphs. In [Appendix D](#), we discuss the role of sequentiality of moves and perfect defense in the results obtained in the paper.

2. THE MODEL

We start with a given network and consider a two-player sequential move game with a defender and an adversary. In the first stage, the defender chooses an allocation of defense resources. In the second stage, given a defended network, the adversary chooses the nodes to attack. Successfully attacked nodes (and their links) are removed from the

⁵For example, in a core–periphery network, all the core nodes are essential separators, while the maximal independent set can include at most one core node and must include peripheral nodes. See [Appendix C](#) for details on this.

network, yielding a residual network. The goal of the defender is to maximize the value of the residual network, while the goal of the adversary is to minimize this value.⁶

Let $N = \{1, \dots, n\}$, with $n \geq 3$, be a finite set of *nodes*. A *link* is a two element subset of N . The set of all possible links over $P \subseteq N$ is $g^P = \{ij : i, j \in P, i \neq j\}$ (where ij is an abbreviation for $\{i, j\}$). A *network* is a set of links. Given the set of nodes $P \subseteq N$, $\mathcal{G}(P) = 2^{g^P}$ is the set of all networks over P . The set $\mathcal{G} = \bigcup_{P \subseteq N} \mathcal{G}(P)$ is the set of all networks that can be formed over any subset of nodes from N . Every network $g \in \mathcal{G}$ has a *value* $\Phi(g)$, associated with it: $\Phi : \mathcal{G} \rightarrow \mathbb{R}$ is called a *value function*.

The set of nodes $X \subseteq N$ chosen by the adversary is called an *attack*. The set $X = \emptyset$ is called the *empty attack*. A *defense* is a set of nodes $\Delta \subseteq N$; node $i \in N$ is defended under Δ if and only if $i \in \Delta$. We assume that the defense is perfect: a protected node cannot be removed by an attack, while any attacked unprotected node is removed with certainty. Given defense Δ and attack X , set $Y = X \setminus \Delta$ will be removed from the network. Removing a set of nodes $Y \subseteq N$ from a network creates a *residual network* $g - Y = \{ij \in g : i, j \in N \setminus Y\}$.

Defense resources are costly: the cost of defending a node is $c_D > 0$. Given network g , the defender's payoff from strategy $\Delta \subseteq N$, when faced with the adversary's strategy $X \subseteq N$, is

$$\Pi^D(\Delta, X; g, c_D) = \Phi(g - (X \setminus \Delta)) - c_D|\Delta|.$$

Attack resources are costly: the cost of attacking a node is given by $c_A > 0$. Given defended network (g, Δ) , the payoff to the adversary from strategy $X \subseteq N$ is

$$\Pi^A(\Delta, X; g, c_A) = -\Phi(g - (X \setminus \Delta)) - c_A|X|. \tag{1}$$

We study the (subgame perfect) equilibria of this game.

Two nodes i and j are connected in network g if there is a sequence of nodes i_0, \dots, i_m such that $i = i_0, j = i_m$, and for all $0 < k \leq m, i_{k-1}i_k \in g$. A *component* of network g is a maximal and nonempty set of nodes $C \subseteq N$ such that any two distinct nodes $i, j \in C$ are connected in g . The set of components of g is denoted by $\mathcal{C}(g)$.

We assume that Φ is component additive. Given network g ,

$$\Phi(g) = \sum_{C \in \mathcal{C}(g)} f(|C|),$$

where f satisfies the following assumption:

ASSUMPTION 1. *We have that $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is strictly increasing, strictly convex, and $f(0) = 0$.*

⁶The sequential move game formulation appears to be appropriate for the large scale and time consuming protection investments discussed in the Introduction. This two-stage model with observability of first-stage actions is consistent with the approach in the large literature on security and networks; see, e.g., Tambe (2012) and Alpcan and Başar (2011).

Remarks on the model

We have assumed sequential moves; this is mainly for exposition. It is possible to show that our main results on characterization of conflict in terms of certain properties of the graph carries over with simultaneous moves. Perfect defense is a more substantial assumption. Smoother models of conflict such as the Tullock contest function would lead to modifications in parts of the main characterization results below. [Appendix D](#) discusses these points in greater detail. Finally, we have assumed that payoffs depend only on the sizes of the networks (or their components): so we abstract from other topological details of the network. This simplification allows us to make progress and should be seen as a first step in the study of network defense.

Since the game is finite and sequential, standard results guarantee the existence of (subgame perfect) equilibria. These equilibria are usually not unique, but generically, equilibrium outcomes are equivalent with respect to player's payoffs, sizes of defense and attack, and the value of residual network. This is the content of the following result.

PROPOSITION 1. *For any network g and costs c_D and c_A , there exists a subgame perfect equilibrium. For generic values of c_A and c_D and generic f , the equilibrium attack and defense size and the payoffs of the players are unique.*

3. THE ANALYSIS

This section develops our main results for the two-person game between the defender and the adversary. Optimal attack focuses on sets of nodes that fragment the network (the separators), while optimal defense targets sets of nodes that block these separators (the transversals). The interest then moves on to the relation between network architecture and the intensity of conflict (the sum of resources allocated to attack and defense) and the prospects of active conflict (when the adversary eliminates some nodes while the defender protects others).

We begin with a study of a simple example that helps illustrate a number of interesting phenomena.

EXAMPLE 1 (Defense and attack on the star). Consider the star network with $n = 4$ and $\{a\}$ as the central node (as in [Figure 1](#)). The value function is $f(x) = x^2$.

As is standard, we solve the game by working backward. For every defended network (g, Δ) we characterize the optimal response of the adversary. We then compare the payoffs to the defender from different profiles, (g, Δ) , and compute the optimal defense strategy. Equilibrium outcomes are summarized in [Figure 2](#).

A number of points are worth noting.

- (i) Observe that removing node a disconnects the network; this node is a separator. Moreover, there is a threshold level of cost of attack such that the adversary either attacks a or does not attack at all when $c_A > 7$. Protecting this node is also central to network defense.




FIGURE 1. Star network ($n = 4$).




FIGURE 2. Equilibrium outcomes: star network ($n = 4$) and $f(x) = x^2$.

- (ii) The intensity of conflict exhibits rich patterns: when the cost of attack is very large there is no threat to the network and no need for defense. If the cost of attack is small, the intensity of conflict hinges on the level of defense costs. When they are low, all nodes are protected and there is no attack (the costs of conflict are nc_D); if they are high, then there is no defense but all nodes are eliminated (the costs of conflict are nc_A). For intermediate cost of attack and defense, both defense and attack are seen in equilibrium.
- (iii) The size of defense may be nonmonotonic in the cost of attack. Fix the cost of defense at $c_D = 3.5$. At a low cost of attack ($c_A < 1$) the defender protects all nodes, in the range $c_A \in (1, 5)$ he protects 0 nodes, in the range $c_A \in (5, 13)$ he protects $\{a\}$, and then in the range $c_A > 13$, he stops all protection activity. Similarly, the size of the attack strategy may be nonmonotonic in the cost of attack. \diamond




FIGURE 3. Essential separators.

The starting point of the general analysis is the nature of optimal attack. Given the convexity in the value function of networks, disconnecting a network is especially damaging. A set $X \subseteq N$ is a separator if $|\mathcal{C}(g)| < |\mathcal{C}(g - X)|$. In other words, a separator is a set of nodes the removal of which strictly increases the number of components in the network. A network will normally possess multiple separators and the adversary should target the most effective ones. A separator $S \subseteq N$ is *essential* for network $g \in \mathcal{G}(N)$, if for every separator $S' \subsetneq S$, $|\mathcal{C}(g - S)| > |\mathcal{C}(g - S')|$, i.e., a strict subset of eliminated nodes would lead to a strictly smaller number of components. The set of all essential separators of a network g is denoted by $\mathcal{E}(g)$. Figure 3 illustrates essential separators in an example.

The second element is the level of costs. As illustrated by Example 1, the network defense problem can be divided into two parts, depending on the cost of attack. Given $x \in \mathbb{N}$, $\Delta f(x) = f(x + 1) - f(x)$ is the marginal increase in the value of a component of size x when a single node is added to it. Under Assumption 1, $\Delta f(x)$ is strictly increasing. It is useful to separate two levels of costs: one, high costs with $c_A > \Delta f(n - 1)$, and two, low costs with $c_A < \Delta f(n - 1)$.

We start with the case of high cost as it brings out some of the main general insights in a straightforward way. Facing a high cost, the adversary must disconnect the network, i.e., choose a separator or not attack the network at all. Clearly, the adversary would never use an essential separator that yields a lower payoff than the empty attack. Given the cost of attack c_A and network g , the set of individually rational separators is $\mathcal{E}(g, c_A) = \{X \in \mathcal{E}(g) : \Phi(g) - \Phi(g - X) \geq c_A |X|\}$.

When the cost of attack is low, it may be profitable for the adversary to use attacks that merely remove nodes from the network, without disconnecting it. A set $R \subseteq N$ is a reducing attack for a network g if there is no $X \subseteq R$ such that X is a separator for g . The set of all reducing attacks for a given network g is denoted by $\mathcal{R}(g)$.

The following lemma characterizes all the possible attacks of the adversary in terms of essential separators and reducing attacks. In addition, it provides a characterization of the attacks that are best responses in the adversary’s subgame.

LEMMA 1. *Fix a connected network g . Let $\Delta \subseteq N$ be a defense selected by the defender in the first stage. Any attack $X \subseteq N$ can be decomposed into two disjoint sets: a set E and a reminder set R such that the following statements hold:*

- (i) *The set E is either empty or $E \in \mathcal{E}(g)$.*
- (ii) *The set R is a reducing attack for $g - E$.*

Moreover, if X is a best response to Δ , then E is either empty or $E \in \mathcal{E}(g, c_A)$.

The first part of the lemma says that any attack of the adversary can be seen as consisting of two phases. In one of the phases the adversary fragments the network by removing a minimal set of nodes needed to obtain the desired components after the attack. This set is an essential separator of the network. In the other phase the adversary reduces the size of the components (but without disconnecting any of them). Thus the notion of essential separator captures exactly the attacks that serve the function of fragmenting the network. The characterization of attacks obtained in the first part of the lemma is useful in understanding the best responses of the adversary. If X is a best response to some strategy of the defender, then applying an essential separator phase of X after the reducing attack phase is applied must be worthwhile. But then, by convexity of f , it must be worthwhile even more to apply the essential separator phase before the reducing attack phase. Therefore the essential separator phase must be individually rational.

We now turn to the equilibrium strategies of the defender. Again, it is instructive to start with the setting where the cost of attack is high. An optimal strategy of the defender should block a subset of individually rational essential separators in the most economical way. Given a family of sets of nodes, \mathcal{H} , and a set of nodes, M , $\mathcal{D}(M, \mathcal{H}) = \{X \in \mathcal{H} : X \cap M \neq \emptyset\}$ are the sets in \mathcal{H} that are blocked (or covered) by M . The set M is called a transversal of \mathcal{H} if $\mathcal{D}(M, \mathcal{H}) = \mathcal{H}$. The set of all transversals of \mathcal{H} is denoted by $\mathcal{T}(\mathcal{H})$. Elements of $\mathcal{T}(\mathcal{H})$ that are minimal with respect to inclusion are called *minimal* transversals of \mathcal{H} . Elements of $\mathcal{T}(\mathcal{H})$ with the smallest size are called *minimum* transversals of \mathcal{H} . Let $\tau(\mathcal{H})$ denote the *transversal number* of \mathcal{H} , i.e., the size of a minimum transversal of \mathcal{H} . Given a family of sets $\mathcal{F} \subseteq \mathcal{H}$, the set M is called a transversal of \mathcal{F} in \mathcal{H} if $\mathcal{D}(M, \mathcal{H}) = \mathcal{F}$. The set of all transversals of \mathcal{F} in \mathcal{H} is denoted by $\mathcal{T}(\mathcal{F}|\mathcal{H})$. Elements of $\mathcal{T}(\mathcal{F}|\mathcal{H})$ with the smallest size are called *minimum* transversals of \mathcal{F} in \mathcal{H} . Let $\tau(\mathcal{F}|\mathcal{H})$ denote the transversal number of \mathcal{F} in \mathcal{H} , i.e., the size of a minimum transversal of \mathcal{F} in \mathcal{H} . Notice that $\tau(\mathcal{F}|\mathcal{H}) \geq \tau(\mathcal{F})$. In other words, avoiding blocking some of the potential attacks of the adversary, and hence strategically exposing some parts of the network, may entail an additional cost. As we show below, strategic exposure may be a part of a rational defense strategy.

Let g be the network in Figure 3. Let $\mathcal{H} = \mathcal{E}(g) = \{\{a\}, \{a, b\}, \{a, c\}\}$ be the set of all essential separators of g and let $\mathcal{F} = \{\{a, b\}, \{a, c\}\}$ be its subset. Figure 4 illustrates the unique minimum (and, at the same time, minimal) transversal of \mathcal{H} , $\{a\}$. The unique minimum transversal of \mathcal{F} in \mathcal{H} is $\{b, c\}$. Thus the most economic way to block exactly the separators from \mathcal{F} out of all the separators from \mathcal{H} is by blocking nodes b and c .




FIGURE 4. Minimum transversal, $\{a\}$, of essential separators $\{\{a\}, \{a, b\}, \{a, c\}\}$.

We provide more examples and a discussion of essential separators and their transversals in some well known families of networks (trees, core–periphery, interlinked stars) in Appendix C.

We are now ready to state our first main result on optimal defense and attack.

PROPOSITION 2. *Consider a connected network $g \in \mathcal{G}(N)$. Let (Δ^*, X^*) be an equilibrium.*

(i) *If $c_A < \Delta f(n - 1)$, then the following statements hold:*

- *The variable $\Delta^* = N$ or Δ^* is a minimal transversal of $\mathcal{D}(\Delta^*, \mathcal{E}(g, c_A))$.*
- *We have $X^*(\Delta) = E \cup R$, where $E \in \mathcal{E}(g, c_A)$ and $R \in \mathcal{R}(g - E)$, with $X^*(\Delta) \cap \Delta = \emptyset$.*

(ii) *If $c_A > \Delta f(n - 1)$, then the following statements hold:*

- *We have $|\Delta^*| \leq \tau(\mathcal{E}(g, c_A))$ and Δ^* is a minimum transversal of $\mathcal{D}(\Delta^*, \mathcal{E}(g, c_A))$ in $\mathcal{E}(g, c_A)$.*
- *We have $X^*(\Delta) = \emptyset$ if $\Delta \in \mathcal{T}(\mathcal{E}(g, c_A))$; $X^*(\Delta) \in \mathcal{E}(g, c_A)$ with $X^*(\Delta) \cap \Delta = \emptyset$, otherwise.*

The proposition brings out the economic trade-offs in the network conflict. Essential separators—that are effective at fragmenting the network—are key to optimal attack and economical transversals that block these separators are key to optimal defense. Moreover, if the defender wishes to go beyond blocking the separator and protect nodes that merely expand the size of a component, then, due to convex character of network value function, it is optimal for him to protect all the nodes in the network.

More formally, optimal defense is defined in terms of the minimal transversal of the appropriate set of essential separators or defense must cover all nodes. If the cost of attack is such that elimination of single nodes is not worthwhile, optimal attack is

bounded from above by the set of essential separators of the network. In this case, optimal defense can never exceed the size of the minimum transversal of the set of individually rational essential separators. If, alternatively, the cost of attack justifies the elimination of single nodes, optimal attack is constituted of nodes that comprise reducing attacks and essential separators. In this case, an interesting feature of optimal defense is that it may be larger than the smallest possible transversal (even when it does not cover all the nodes).

We now briefly describe the arguments underlying the proof. By [Lemma 1](#), we know that any attack may be decomposed into two disjoint parts that comprise an essential separator and a reducing attack.

In the range of costs covered by part (ii), the adversary will not use reducing attacks. So, an optimal attack must be either empty or an individually rational essential separator. Next consider the optimal defense strategy, Δ^* . Clearly, Δ^* cannot be larger than the size of the minimum transversal of $\mathcal{E}(g, c_A)$, as that would be wasteful for the defender. If $|\Delta^*| = \tau(\mathcal{E}(g, c_A))$, then Δ^* must be a minimum transversal of $\mathcal{E}(g, c_A)$; choosing a defense other than a minimum transversal would simply lower payoffs. If $|\Delta^*| < \tau(\mathcal{E}(g, c_A))$, then Δ^* is a minimum transversal of $\mathcal{D}(\Delta^*, \mathcal{E}(g, c_A))$ in $\mathcal{E}(g, c_A)$.

We turn next to part (i) of [Proposition 2](#). The proof proceeds by showing that a defense that exceeds a minimal transversal (of covered essential separators) must include some node that is being protected purely to prevent it from removal. Hence the role of such a defense is to ensure the size of the component. This must mean that, in the absence of defense, the node would be eliminated in the subsequent optimal attack. We then exploit the convexity of f and the linearity of costs of defense and attack to establish that the adversary must find it optimal to eliminate all other unprotected nodes in the surviving component. Extrapolating from this, we establish that this must apply to all essential separators and then, by convexity, to single nodes in those components as well. In other words, if the defender finds it optimal to go beyond a minimal transversal of blocked essential separators, then he must protect all nodes.

We now consider the general comparative statics with respect to the costs and the network. It is worth noting some patterns in [Example 1](#) above. [Figure 2](#) suggests that defense size is falling in defense costs and is nonmonotonic in attack costs. The attack size is nonmonotonic in both attack cost and defense cost. These patterns are true more generally. They have payoff implications. The following result summarizes our analysis.

PROPOSITION 3. *The equilibrium comparative statics are as follows.*

- (i) *The size of defense and the defender's payoff are both decreasing in the cost of defense. The defender's payoff increases in the cost of attack. However, depending on the costs and the network, the size of defense may increase or decrease when the cost of attack increases.*
- (ii) *Depending on the costs and the network, the size of attack and adversary's payoff may increase or decrease when the cost of attack increases. The adversary's payoff increases in the cost of defense. However, depending on the costs and the network, the size of attack may increase or decrease when the cost of defense increases.*




FIGURE 5. Network where a rise in the cost of defense reduces the size of attack.

(iii) Depending on the costs and the network, adding links may increase or decrease the size of the optimal defense as well as the defender’s payoff.

We note that the effect of defense cost on the size of attack may be nonmonotonic. This is because with a higher cost of defense, the defender may uncover some essential separators to which the adversary could switch. Their size might be smaller or larger than the size of separators chosen by the adversary under the lower cost of defense. As an example, consider the network g in Figure 5 and suppose that $f(x) = x^2$, $c_A \in (31, 54)$, and $c_D \in (108, 121)$. Under these parameters, in every equilibrium the defender defends node a and the adversary responds with essential separator $\{b, c\}$. When the cost of defense rises to 122, equilibrium defense of the defender is \emptyset to which the adversary responds with essential separator $\{a\}$. Alternatively, Example 1 illustrates that the size of attack might rise when the cost of defense is rising (cf. the case of $c_A \in (7, 13)$ in Figure 2). Despite this nonmonotonic behavior of equilibrium attack size, the payoff to the adversary increases when the cost of defense rises. A similar observation also holds for the effect of attack cost on defense size and on payoffs.

An increase in attack cost has nonmonotonic effects on attack size and the adversary’s payoff. This is illustrated by Example 1, e.g., when the cost of defense is in the range $(3.25, 4)$. The reason for these nonmonotonicities is as follows. When the cost of attack rises, some of the attacks stop being individually rational. This creates an opportunity for the defender to reduce defense, possibly at the expense of some value of the network. This, in turn, allows the adversary to execute attacks that were blocked when the cost of attack was lower. In the example, when $c_A \in (0, 1)$, it is individually rational for the adversary to remove any unprotected node. Therefore, with $c_D \in (3.25, 4)$, the defender defends all the nodes. When $c_A \in (1, 5)$, it is not individually rational for the adversary to remove single unprotected nodes. With the costs of defense in $(3.25, 4)$, the defender prefers to leave the network undefended and loose the central node, saving on the cost of defense and losing some value of the network. Such an attack is better for the adversary than not removing any node. The size of attack rises from 0 to 1 and the payoff of the adversary rises from -16 to $-3 - c_A \in (-9, -4)$. When $c_A > 7$, the size of attack falls back to 0 and the payoff to the adversary falls back to -16 .

Finally, consider the effects of adding links. A first conjecture would be that adding links should always be good for the defender, as it creates more routes for connection and this should make the network easier to defend. The next example shows that this intuition is false: a denser network may induce a larger optimal defense with lower defender payoffs!




FIGURE 6. Example 2. (a) Original network. (b) Network with added link.

EXAMPLE 2 (Adding links may increase defense size and lower defender payoffs). We consider the network given in Figure 6. Suppose that payoff from a component of size x is $f(x) = x^2$.

Assume the cost of attack $c_A \in (23, 31)$ and the cost of defense $c_D \in (43, 85)$. The unique equilibrium outcome is $\Delta^* = \{c\}$, $X^* = \{d\}$. The equilibrium payoff to the defender is $101 - c_D$.

Now consider a network $g' = g \cup \{ef\}$, with a link added between the nodes e and f . With this additional link, the separator $\{d\}$ is replaced by separator $\{d, e\}$. Suppose that the cost of defense is $c_D \in (43, 62)$. Observe that with defense $\Delta^* = \{c\}$, there exists an attack $\{d, e\}$ that is optimal for the adversary and yields only $82 - c_D$ to the defender. Thus the addition of a link, and retaining the same defense, may actually lower the defender's payoffs.

In the new network g' , the unique equilibrium outcome is $\Delta^* = \{d, e\}$ and $X^* = \emptyset$. The equilibrium payoff to the defender is $144 - 2c_D < 101 - c_D$. So, *the optimal defense size increases and the defender's payoff falls as the network becomes denser*.

Alternatively, it is clear that as we keep adding links and arrive at the complete network, the optimal attack is empty (as $c_A > 23$) and so optimal defense is also the empty set. The defender's payoff is 144, which is the maximal attainable. Thus the effects of adding links are nonmonotonic. \diamond

This nonmonotonicity is not an artefact of the specifics of the network and the costs of attack and defense. It reflects a general feature of conflict in networks. To see this consider the case of the complete network. The first thought would be that a network that contains the most connections is the hardest to disrupt and always leads to the best outcomes for the defender. This is not true. The following example clarifies this point.

EXAMPLE 3 (Complete network vs. core–periphery network). Suppose that n is large and that the cost of attack satisfies

$$f(n - 2) - f(n - 3) < c_A < f(n - 1) - f(n - 2).$$

With this cost of attack, the adversary removes two nodes from the complete network over n nodes, one node from the complete network containing $n - 1$ nodes, and does

not remove any nodes from the complete network containing $n - 2$ or less nodes. Finally, suppose that the cost of defense satisfies

$$\frac{f(n) - f(n - 2) - f(1)}{n} < c_D < \frac{f(n) - f(n - 2)}{n}.$$

With this cost of defense the defender protects all the nodes in a complete network with n nodes, because $f(n) - nc_D > f(n - 2)$ (and we know that in a complete network the defender either protects all or no nodes, in equilibrium).

Now consider a network with $n - 1$ nodes in a clique with one node linked to a single element of the core (let us call it i). This is a type of core-periphery network. If such a network is not protected, the adversary will remove node i only, disconnecting the network into a clique of size $n - 2$ and a single isolated node. Now, we know that the defender is either inactive, protects i , or protects all the nodes in equilibrium. With the above cost of defense, the defender is inactive. First, note that $f(n) - nc_D < f(n - 2) + f(1)$, so protecting everything is worse than being inactive. It can be checked that protecting i is worse, because in response the adversary would remove two nodes from the core of the network.

Thus in the core-periphery network the equilibrium payoff to the defender is

$$f(n - 2) + f(1) > f(n) - nc_D.$$

So it is better than the complete network. ◇

This example illustrates the attractiveness of the queen sacrifice strategy: it is better to leave i unprotected because there is greater loss in value if it is protected! The idea of queen sacrifice and the suboptimality of the complete network will resurface in other contexts below.

3.1 Networks and conflict

This section examines the relation between the network architecture and the nature of conflict more closely. We define the *intensity of conflict* as the sum of expenditures of defense and attack. Our analysis shows that for given costs of conflict, differences in network structure can lead to very large differences in conflict.

Proposition 1 tells us that the size of equilibrium attack and defense are generically unique. We start by defining the minimum intensity of conflict for given costs of attack and defense. Define minimal costs of conflict for given costs and f as

$$CC(c_A, c_D, f) = \min_{g \in \mathcal{G}(N)} c_D |\Delta^*(g, c_A, c_D, f)| + c_A |X^*(g, c_A, c_D, f)|.$$

Example 1 illustrates some of the forces at work. Observe that when the cost of attack is very large, $c_A > 13 = f(n) - (n - 1)f(1)$, no attack is profitable, and, anticipating this, the defender abstains from defense. The intensity of conflict is 0. This lack of conflict for large costs of attack is independent of the architecture of the network.




FIGURE 7. Windmill graphs (h_n^m): $n = 13, m = 6, 4, 3$.

Turning to the lower cost of attack, an inspection of Figure 1 in Example 1 tells us that the intensity of conflict also depends on the cost of defense. It will be useful to define a special class of networks, windmill graphs. These graphs are denoted by h_n^m , where $n \geq 2$ and $m \in \{1, \dots, n - 1\}$. There is one critical node that, when removed, disconnects the network. The remaining nodes are partitioned into cliques of size m and, possibly, one clique of smaller size (this implies that there are $\lceil (n - 1)/m \rceil$ such cliques). Every member of a clique is connected to the critical node. We now define a key cost threshold for defense that equates the payoff from full defense with the payoff from an unprotected h_n^m network:

$$c(m, n) = \frac{f(n) - \lfloor \frac{n-1}{m} \rfloor f(m) - f((n - 1) \bmod m)}{n}.$$

Figure 7 illustrates windmill graphs.

We are now ready to provide a general characterization of minimal conflict levels.

PROPOSITION 4. (i) If $c_A > f(n) - (n - 1)f(1)$, then $CC(c_A, c_D, f) = 0$. It is attained on any connected network.

(ii) If $c_A \in (\Delta f(n - 1), f(n) - (n - 1)f(1))$, then $CC(c_A, c_D, f) = 0$. It is attained on any connected network g with $\mathcal{E}(g, c_A) = 0$.

(iii) If $c_A \in (\Delta f(m - 1), \Delta f(m))$ with $m \in \{1, \dots, n - 1\}$, one of the following statements holds:

(iv) If $c_D > c(m, n)$, then $CC(c_A, c_D, f) = c_A$. It is attained on a windmill network, h_n^m .

(v) If $c_D < c(m, n)$ with $m \in \{1, \dots, n - 1\}$, then $CC(c_A, c_D, f) = nc_D$. It is attained on any connected network.

In case (ii), when the cost of attack is high, $c_A > \Delta f(n - 1)$, the minimal costs of conflict are 0, as it is not profitable for the adversary to attack any network with $\mathcal{E}(g, c_A) = \emptyset$. Such networks include the complete network, as well as networks that are robust to node

removal in the sense that they require a large number of nodes to be removed to get disconnected. More generally, for any integer $t \geq 1$, a network is *t-connected* if it can be disconnected by removing t nodes and cannot be disconnected by removing less than t nodes. Any t -connected network with $t \geq (f(n) - nf(1))/(c_A - f(1))$ has empty $\mathcal{E}(g, c_A)$. Menger (1927) provides a characterization of such networks: a network is at least t -connected if and only if any two nodes that are not neighbors are connected through at least t node independent paths.⁷ Thus such networks have many redundant connections between nodes.

The last case, with lower attack costs $c_A < \Delta f(n - 1)$, is much richer. Suppose that $c_A \in (\Delta f(m - 1), \Delta f(m))$, where $m \in \{1, \dots, n - 1\}$. Now it is profitable to the adversary to attack any undefended node in a component of size greater than m . Hence, the lower bound on costs of conflict is $\min(c_A, nc_D)$. If the cost of defense is sufficiently low, $c_D < c(m, n)$, then complete defense is better than any other defense and the minimal costs of conflict are nc_D . If $c_D > c(m, n)$, then complete defense has higher costs as compared to the outcome with no defense and one attacked node. This leads to total costs of conflict of c_A . To sustain an equilibrium with such costs of conflict, we need a network that has a separator of size 1 and that all components in the residual network have size at most m . The windmill graph possesses exactly this characteristic. This motivates the windmill network: for $m \in \{1, \dots, n - 2\}$, the windmill network h_n^m has such an equilibrium and yields the minimal costs of conflict, c_A .

We now turn to the role of networks in shaping the intensity of conflict. Proposition 4 tells us that network architecture matters only if the costs are as in cases (ii) or (iii).

Consider case (ii). Proposition 4 tells us that $CC(c_A, c_D, f) = 0$ in this range. To see the impact of network architecture, consider a star network. If $c_D < f(n) - (n - 1)f(1)$, then in equilibrium the defender protects the center of the star and the costs of conflict are c_D . Alternatively, if $c_D > f(n) - (n - 1)f(1)$, then in equilibrium the defender chooses the empty defense, the adversary attacks the center of the star, and the costs of conflict are c_A . So, when the costs of attack and defense reach their upper bound, the difference in the costs of conflict between the star network and the minimal attainable is $f(n) - (n - 1)f(1)$. It is easy to see that this can grow without bound as n gets large.

Next consider case (iii), with $m \in \{1, \dots, n - 2\}$. Proposition 4 tells us that the minimum conflict, attained on network h_n^m (for example) is c_A . Suppose $c_D \in (c(m, n), (f(n) - f(m))/n)$ and consider a complete network. The unique equilibrium outcome is full protection and so the costs of conflict are nc_D . When the cost of defense reaches its upper bound and the cost of attack reaches its lower bound, the difference in costs between this minimum and the complete network reaches $f(n) + f(m - 1) - 2f(m)$, which is maximal, $f(n) - 2f(1)$, for $m = 1$. Again, the network architecture can have very large effects on the intensity of conflict.

Active conflict In Proposition 4, minimal conflict is associated with a single active player. An inspection of Figure 2, in Example 1 above, shows us that both players can

⁷Two paths are *node independent* if the only nodes they have in common are the starting and the ending nodes.

be active in equilibrium. This motivates the study of circumstances under which we should expect to see active conflict. [Example 1](#) draws attention to the role of costs: neither the attack nor the defense costs can be too high. Here we briefly discuss the role of the network architecture and the network value function.

We start with an observation that draws upon [Proposition 2](#): for active conflict to arise there must exist an individually rational essential separator. If such a separator does not exist, then convexity of function f together with linearity of costs implies that either none or all nodes are defended. In particular, if g is a complete network, then for all costs and all functions f (satisfying our assumptions), there is no equilibrium with active conflict.

Are there any other (connected) networks with the same property as complete networks? If the marginal value of f is growing sufficiently fast, then no active conflict is possible. Let f satisfy the property, for $x \geq 0$,

$$\Delta f(x) > xf(x), \tag{2}$$

where $\Delta f(x) = f(x + 1) - f(x)$.

The property is satisfied by functions $f(x) = (x + 1)! - 1$ and $(x + 1)^x - 1$, for example. Marginal value in these functions grows so rapidly that adding a single node to a component of size m increases its value more than m times. In effect, the returns from protecting $m < n$ nodes are smaller than average returns from protecting additional $m - n$ nodes. Thus if the defender prefers protecting the first m nodes to no protection, he is even more willing to protect the whole network. Formally, let

$$\Phi^*(m; g, c_A) = \max_{\Delta \subseteq N, |\Delta| \leq m} \min_{X \in \text{BR}(\Delta; g, c_A)} \Phi(g - X(\Delta) \setminus \Delta)$$

be a function that gives the maximum value of the residual network that can be attained from network g when up to m units of defense are used and the cost of attack is c_A ($\text{BR}(\Delta; g, c_A)$ denotes the set of best responses of the adversary to Δ , given g and c_A). Suppose that there is an equilibrium, (Δ^*, X^*) , featuring active conflict. Let $|\Delta^*| = m$. Since there is active conflict, so $1 \leq m \leq n - 1$ and $|X^*(\Delta^*)| \geq 1$. Since Δ^* is better than \emptyset , so $c_D \leq (\Phi^*(m; g, c_A) - \Phi^*(0; g, c_A))/m \leq f(n - 1)$. Alternatively, since Δ^* is better than N , so $c_D \geq (f(n) - \Phi^*(m; g, c_A))/(n - m) \geq (f(n) - f(n - 1))/(n - 1)$. Combining both the inequalities we get $f(n) \leq nf(n - 1)$, which contradicts (2).

4. DECENTRALIZED DEFENSE

In many applications, security decisions are made at the individual node level. This section studies decentralized security choices in a network that is under attack. We begin by showing that the equilibrium choices of the nodes and the adversary can be characterized in terms of transversals and separators of the underlying network. We then show that the welfare gap between decentralized equilibrium and first best outcomes is unbounded: interestingly, individual choice may lead to too little and to too much protection, relative to the choice of a single (centralized) defender.

We consider a two-stage game. In the first stage, each of the nodes in the network decides whether to protect itself or to stay unprotected. These choices are observed by the adversary who then chooses the nodes to attack.

Let $N = \{1, 2, \dots, n\}$, where $n \geq 3$ is the set of players, and let $S_i = \{0, 1\}$ denote the strategy set of node $i \in N$. Here $s_i = 1$ means that the node chooses to defend itself and $s_i = 0$ refers to the case of no defense. These choices are made simultaneously. There is a one-to-one correspondence between a strategy profile of the nodes, $\mathbf{s} \in \{0, 1\}^N$, and the resulting set of defended nodes $\Delta \subseteq N$. So we will use Δ to refer to the strategy profile of the nodes in the first stage.

In the second stage the adversary observes the defended network (g, Δ) and chooses an attack $X \subseteq N$, which leads to a residual network $g - (X \setminus \Delta)$. The payoff to the adversary remains as in the case of the centralized defense and is defined in (1). The payoff to a node depends on whether the node is removed by the attack. A removed node receives payoff 0. Each of the surviving nodes receives an equal share of the value of its component in the residual network,

$$\Pi^i(\Delta, X; g, c_D) = \begin{cases} 0 & \text{if } i \in X \setminus \Delta \\ \frac{f(C(i))}{|C(i)|} - s_i c_D & \text{otherwise,} \end{cases}$$

where $C(i)$ is the component in the residual network $g - (X \setminus \Delta)$ that contains i .

This completes the description of the *decentralized defense game*. We study the subgame perfect equilibria of this game, restricting attention to those without active conflict.

Let us solve the game starting from the second stage. As in the two-player game, the adversary chooses either the empty attack or an attack that is a combination of an essential separator and a reducing attack. If the cost of attack is low and there is no active conflict, then either the adversary removes all the nodes or all nodes are protected. In any other outcome the adversary must remove at least one node. If the cost of attack is high and there is no active conflict, then either none of the nodes protects or, anticipating the strategy of the adversary, the nodes choose a defense configuration that blocks all the individually rational essential separators. Therefore, in equilibrium, they must choose a minimal transversal of $\mathcal{E}(g, c_A)$. We build on these observations to provide the following characterization of equilibria with no active conflict in the decentralized defense game.⁸

PROPOSITION 5. *Consider a connected network $g \in \mathcal{G}(N)$. Let Δ^* be the equilibrium defense.*

- (i) *If $c_D > f(n)/n$, then $\Delta^* = \emptyset$ is the unique equilibrium defense.*

⁸We concentrate on equilibria with no active conflict, because, on one hand, it allows for providing a clean characterization and, on the other hand, it provides a sufficiently rich platform for discussing the sources of inefficiencies when defense decisions are decentralized. All other equilibria in decentralized defense game could be characterized in the same spirit as the characterization provided in Proposition 2 for the centralized defense game.

(ii) If $c_D \leq f(n)/n$, one of the following statements holds.

(a) If $c_A < f(n) - f(n - 1)$, then $\Delta^* = N$ is an equilibrium defense.

(b) If $c_A > f(n) - f(n - 1)$, then any minimal transversal of $\mathcal{E}(g, c_A)$ is an equilibrium defense.

The equilibrium strategy of the adversary is as in Proposition 2.

We now turn to discussing inefficiencies that may arise due to decentralized protection, as well as their sources. We compare the aggregate welfare of the nodes in the equilibrium of the two-player game with the aggregate welfare in the decentralized defense game. Let $\Pi^{D^*}(g, c_A, c_D)$ denote the equilibrium payoff in the two-player game on network g with cost of defense c_D and cost of attack c_A . Aggregate welfare in the two-player game, starting from network g , and costs c_A and c_D , are defined as

$$W^F(g, c_A, c_D) = \Pi^{D^*}(g, c_A, c_D).$$

Aggregate welfare under defense profile Δ and attack X , of the $n + 1$ -player game starting from network g , and given cost of defense c_D , is defined as

$$W^D(\Delta, X; g, c_D) = \sum_{i \in N} \Pi^i(\Delta, X; g, c_D).$$

Following Koutsoupias and Papadimitriou (1999), we study the cost of decentralization in terms of the *price of anarchy* (PoA): the ratio of welfare in the two-player game to the welfare in the worst equilibrium of the decentralized defense game. Let $E(g, c_A, c_D)$ denote the set of equilibria of the $n + 1$ -player game on network g with cost of attack c_A and cost of defense c_D . Let

$$\text{PoA} = \max_{g, c_A, c_D} \left(\frac{W^F(g, c_A, c_D)}{\min_{(\Delta, X) \in E(g, c_A, c_D)} W^D(\Delta, X(\Delta); g, c_D)} \right).$$

Our analysis highlights externalities and points to sources of inefficiency in decentralized defense. The first source is the familiar one of positive externalities: an individual’s protection decision creates benefits for other nodes, which she does not take into account. Consider a star network and suppose that cost of attack is high, $c_A > f(n) - f(n - 1)$, and $c_D \in (f(n)/n, f(n))$. In the equilibrium of the two-player game, the aggregate welfare is $f(n) - c_D$. However, in the equilibrium of the decentralized game, the central player does not find it profitable to defend itself, as $c_D > f(n)/n$. So aggregate welfare in equilibrium of the $n + 1$ -player game is 0. The ratio of the two is unbounded for $c_D \in (f(n)/n, f(n))$.

Protection choices exhibit a threshold property: for a node to find it profitable to protect it is necessary that other nodes belonging to the same minimal transversal protect. Thus protection decisions are strategic complements. This can generate coordination failures, resulting in large welfare losses. To see this, consider a tree with two hubs each of which is linked to $(n - 2)/2$ distinct nodes. Suppose that

$$f(n) - f(n - 1) < c_A < f\left(\frac{n}{2}\right) - \frac{(n - 2)f(1)}{2}$$




FIGURE 8. Network with essential separators of size 2 having two minimal transversals: one of size 1 and one of size 5.

so the adversary will only attack hub nodes. If $2f(n/2)/n < c_D < f(n)/n$, then the first best outcome is to defend the two hubs. One hub protecting itself gives incentives to the other hub to protect: two protected hubs is an equilibrium outcome. However, a hub node does not have unilateral incentives to protect: zero protection is also an equilibrium outcome. In this equilibrium the aggregate payoffs are $(n - 2)f(1)$ as compared to first best outcome of $f(n) - 2c_D$. The cost of decentralization can be unbounded.

Third, at the local level, the game is clearly one of strategic substitutes. A node in a separator has incentives to protect only if no other node in the separator protects itself. Like public good games on networks (cf. Bramoullé and Kranton 2007), the network protection game therefore displays multiple equilibria. This can generate very large efficiency losses. As an example consider network g depicted in Figure 8.

Suppose that $f(x) = x^2$, $c_A \in (21, 28)$, and $c_D < 11$. Since the cost of attack is high, the adversary will not remove a node without disconnecting the network. The set of individually rational essential separators is a combination of sets as depicted in Figure 8. Notice that the minimum transversal of $\mathcal{E}(g, c_A)$ is the node belonging to each of the separators, while the largest minimal transversal consists of one distinct node from each of the two element separators. Hence the modified PoA in this case is $|\mathcal{E}(g, c_A)|$ and as the example in Figure 8 suggests, it is possible to have a graph g such that $|\mathcal{E}(g, c_A)| \geq (n - 1)/2$. Again, the cost of decentralization is unbounded.

The idea that personal security exhibits positive externalities is well known in the economic epidemiology literature (and has been noted in the recent research in this area; see, e.g., Acemoglu et al. (2013), Cerdeiro et al. (2014), Zawadowski (2013). Moreover, in the standard disease setting security choices are strategic substitutes. Our model departs from this standard setting in two important ways: one, we have an intelligent adversary, and two, agents in our model care about the size of the component (not just about survival). This means that security choices exhibit features of both complements and substitutes. In addition due to the role of size effects, security choices can exhibit large coordination failures. These features of the model distinguish it from the existing literature and call for new methods of analysis and yield fresh insights.

5. CONCLUDING REMARKS

Infrastructure networks are a key feature of an economy. These networks face a variety of threats ranging from natural disasters to intelligent attacks. This paper develops a strategic model of defense and attack in networks.

We provide a characterization of equilibrium attack and defense in terms of two classical concepts in graph theory: separators and transversals. We show that the intensity of conflict (the resources spent on attack and defense) and the possibility of active conflict (when both adversary and defender target nodes for action) are both intimately related to the architecture of the network. Finally, we show that the welfare costs of decentralized defense can be very large.

We have assumed that the defender moves first and is followed by the attacker, and that the defense is perfect: it would be more natural to allow for outcomes of conflict to vary with resources of attack and defense allocated to a node. [Appendix D](#) presents a preliminary analysis of models where we relax these assumptions. A general analysis remains an important problem for future research.

Finally, we have assumed that payoffs depend only on the sizes of the networks (or their components). In future work, it would be important to study a model where payoffs depend on the details of the architecture of the components.

APPENDIX A: PROOFS

We start with proving [Proposition 1](#) that states generic equivalence of equilibrium outcomes of the defender–adversary game in terms of payoffs, size of defense, and size of attack. We start with the following auxiliary lemmata.

LEMMA 2. *Let g be a network over set of nodes N and let $\Delta \subseteq N$ be a set of defended nodes. Generically, for any best responses X^* and X^{**} to defense Δ , $\Phi(g - X^*) = \Phi(g - X^{**})$ and $|X^*| = |X^{**}|$.*

PROOF. Let g be a network and let Δ be a defense, as stated in the lemma. Let X^* and X^{**} be best responses to (g, Δ) . Then we have

$$-\Phi(g - X^*) - |X^*|c_D = -\Phi(g - X^{**}) - |X^{**}|c_D.$$

If $|X^*| = |X^{**}|$, then it follows that $\Phi(g - X^*) = \Phi(g - X^{**})$ and we are done. Otherwise, the equality is equivalent to

$$c_D = \frac{\Phi(g - X^*) - \Phi(g - X^{**})}{|X^{**}| - |X^*|}.$$

The set of values on the right-hand side of the equality is finite (there are at most $2^{n+1} - 1$ values there). Hence the equality can be satisfied for a finite number of values of $c_D \in \mathbb{R}_{++}$ and it is not satisfied for almost any value of $c_D \in \mathbb{R}_{++}$. This completes the proof. □

LEMMA 3. Let g be a network over the set of nodes N . Generically, for any two equilibria (Δ^*, X^*) and (Δ^{**}, X^{**}) , $\Phi(g - X^*(\Delta^*)) = \Phi(g - X^{**}(\Delta^{**}))$ and $|\Delta^*| = |\Delta^{**}|$.

PROOF. Let $g, \Delta^*, \Delta^{**}, X^*$, and X^{**} be as stated in the lemma. Since Δ^* is a best response to X^* , so

$$\Phi(g - X^*(\Delta^*)) - |\Delta^*|_{c_D} \geq \Phi(g - X^*(\Delta^{**})) - |\Delta^{**}|_{c_D}, \tag{3}$$

and since Δ^{**} is a best response to X^{**} , so

$$\Phi(g - X^{**}(\Delta^{**})) - |\Delta^{**}|_{c_D} \geq \Phi(g - X^{**}(\Delta^*)) - |\Delta^*|_{c_D}. \tag{4}$$

By Lemma 2, generically, $\Phi(g - X^{**}(\Delta^*)) = \Phi(g - X^*(\Delta^*))$ (as both $X^{**}(\Delta^*)$ and $X^*(\Delta^*)$ are best responses to Δ^*). This together with (3) and (4) implies

$$\Phi(g - X^{**}(\Delta^{**})) - |\Delta^{**}|_{c_D} \geq \Phi(g - X^*(\Delta^*)) - |\Delta^*|_{c_D}.$$

Similarly, by Lemma 2, generically, $\Phi(g - X^*(\Delta^*)) = \Phi(g - X^{**}(\Delta^*))$. This together with (3) and (4) implies

$$\Phi(g - X^{**}(\Delta^{**})) - |\Delta^{**}|_{c_D} = \Phi(g - X^*(\Delta^*)) - |\Delta^*|_{c_D}. \tag{5}$$

If $|\Delta^*| = |\Delta^{**}|$, then $\Phi(g - X^*(\Delta^*)) = \Phi(g - X^{**}(\Delta^{**}))$ and we are done. Otherwise, (5) can be rewritten as

$$c_D = \frac{\Phi(g - X^*(\Delta^*)) - \Phi(g - X^{**}(\Delta^{**}))}{|\Delta^*| - |\Delta^{**}|}.$$

Since the number of values on the right-hand side is finite, for almost every value of $c_D \in \mathbb{R}_{++}$ this equality is not satisfied. Hence, generically, $|\Delta^*| = |\Delta^{**}|$ and $\Phi(g - X^*(\Delta^*)) = \Phi(g - X^{**}(\Delta^{**}))$. □

LEMMA 4. Let g be a network over set of nodes N and let $X, Y \subseteq N$ be two attacks such that $|X| \neq |Y|$. Generically, $\Phi(g - X) \neq \Phi(g - Y)$.

PROOF. Let g, X , and Y be as stated in the lemma. Suppose that $\Phi(g - X) = \Phi(g - Y)$. This equality can be rewritten as

$$\sum_{C \in \mathcal{C}(g-X)} f(|C|) = \sum_{C \in \mathcal{C}(g-Y)} f(|C|).$$

Since $X \neq Y$ so there exists $s > 0$ such that $g - X$ has a component of size s and $g - Y$ has not or $g - Y$ has a component of such a size and $g - X$ has not. Suppose that $\Phi(g - X) = \Phi(g - Y)$. Hence the equality above reduces to

$$f(s_1) + \dots + f(s_p) = f(z_1) + \dots + f(z_q), \tag{6}$$

where s_1, \dots, s_p and z_1, \dots, z_q are sizes of components such that $\{s_1, \dots, s_p\} \cap \{z_1, \dots, z_q\} = \emptyset$. Equation (6) puts very strict constraints on function f and perturbing it

slightly (within the set of functions satisfying **Assumption 1**) destroys the equality. Thus $\Phi(g - X) = \Phi(g - Y)$ for $|X| \neq |Y|$ is a nongeneric property of f . \square

With **Lemmas 2, 3, and 4** in hand, we are ready to prove **Proposition 1**.

PROOF OF PROPOSITION 1. Generic equivalence of defense size and of payoff to the defender follow directly from **Lemma 3**. Consider equivalence of attack size and of payoff to the adversary. By **Lemmata 3 and 4**, generically $\Phi(g - X^*(\Delta^*)) = \Phi(g - X^{**}(\Delta^{**}))$ and $|X^*(\Delta^*)| = |X^{**}(\Delta^{**})|$. Thus the points follow as well. \square

Proofs of **Lemma 1** and **Proposition 2** exploit some properties of graphs. The first step is to establish these properties. **Lemma 5** characterizes the essential separators as those separators that are “thin”: every node of such separators is a neighbor of at least two components of the residual network. Given a set of nodes $X \subseteq N$ and a network g over N , $\partial_g(X) = \{k \in N \setminus X : \text{there is } j \in X \text{ such that } jk \in g\}$ is the *neighborhood of X in g* . If X is a singleton, that is, $X = \{j\}$, then we will write $\partial_g(j)$ instead of $\partial_g(\{j\})$ ($\partial_g(j)$ is the set of neighbors of j in g). We will drop the subscript g in the notation if network g is clear from the context.

LEMMA 5. *Let $g \in \mathcal{G}(N)$ be a network over a set of nodes N . A set $X \subseteq N$ is an essential separator if and only if $X \neq \emptyset$ and for every $i \in X$ there exist two distinct components $C_1, C_2 \in \mathcal{C}(g - X)$, $C_1 \neq C_2$, such that $\partial_{g-X}(i) \cap C_1 \neq \emptyset$ and $\partial_{g-X}(i) \cap C_2 \neq \emptyset$.*

PROOF. Let $g \in \mathcal{G}(N)$ be a network over a set of nodes N and let $X \subseteq N$.

The necessary part. Assume that X is an essential separator. Since X is a separator, so $X \neq \emptyset$. Assume, to the contrary, that there exists $i \in X$ such that there is at most one component $C \in \mathcal{C}(g - X)$ such that $\partial_{g-X}(i) \cap C \neq \emptyset$. Suppose first there is no such component. Then the attack $X' = X \setminus \{i\}$ results in the set of components $\mathcal{C}(g - X') = \mathcal{C}(g - X) \cup \{\{i\}\}$, larger than $\mathcal{C}(g - X)$, which contradicts the assumption that X is essential. Second, suppose that there is exactly one component $C \in \mathcal{C}(g - X)$ such that $\partial_{g-X}(i) \cap C \neq \emptyset$. Taking attack X' , as before, leads to a residual network with set of components $\mathcal{C}(g - X') = (\mathcal{C}(g - X) \setminus \{C\}) \cup \{C \cup \{i\}\}$, which has the same cardinality as $\mathcal{C}(g - X)$. Therefore X is not essential, a contradiction.

Sufficiency part. Assume that $X \neq \emptyset$, and for every $i \in X$, there exist two distinct components $C_1, C_2 \in \mathcal{C}(g - X)$ such that $C_1 \cap \partial_{g-X}(i) \neq \emptyset$ and $C_2 \cap \partial_{g-X}(i) \neq \emptyset$. Then there exist two nodes, $j_1 \in C_1 \cap \partial_{g-X}(i)$ and $j_2 \in C_2 \cap \partial_{g-X}(i)$, that are connected in g and not connected in $g - X$. Hence X is a separator and we have to show that it is essential. Suppose $X' \subsetneq X$, so there is some i such that $i \in X$ but $i \notin X'$. Given the definition of $i \in X$ it follows that $|\mathcal{C}(X')| \leq |\mathcal{C}(X)| - 1$. Since X' was arbitrary, the claim is established. \square

We now develop a characterization of optimal attack strategies in terms of essential (individually rational) separators and reducing attacks.

PROOF OF LEMMA 1. The proof of the first part is by induction on the number of nodes in X that violate the condition from **Lemma 5**. For the induction basis consider the set of all

$X \subseteq N$ for which there are no nodes that violate the condition. Then, by Lemma 5, X is essential and so the remainder is \emptyset and $E = X$ (in particular, it may be that $E = X = \emptyset$). The claim holds.

For the induction step, take any $X \subseteq N$ for which there are exactly m nodes that violate the condition from Lemma 5. Suppose that the claim holds for any $Y \subseteq N$ for which there are $l < m$ nodes that violate the condition. Let $i \in X$ be a node that violates the condition and let $Y = X \setminus \{i\}$. Since the condition is violated for $i \in X$, so $g - Y$ either contains one component more than $g - X$ (namely component $\{i\}$) or it has the same number of components with one component C in $g - X$ replaced with $C \cup \{i\}$ in $g - Y$. Hence the condition is violated for $l < m$ nodes from Y in $g - Y$. Thus, by the induction hypothesis, Y can be decomposed into two disjoint sets E and R as claimed. Since, as we argued above, adding i to Y does not increase the number of components in the residual network, so $R \cup \{i\}$ does not contain a separator of $g - E$ and so the decomposition of X into E and $R \cup \{i\}$ satisfies the conditions from the claim. Thus points (i) and (ii) are shown.

Now we show that if g is connected and X is a best response to some defense $\Delta \subseteq N$, then either $E = \emptyset$ or $E \in \mathcal{E}(g, c_A)$.

We show first, for any attack X and any decomposition of X into two disjoint sets E and R satisfying points (i) and (ii), that

$$\Phi(g - E) - \Phi(g - X) \leq \Phi(g) - \Phi(g - R). \tag{7}$$

We use induction on R . For the induction basis, let $R = \emptyset$. Then (7) trivially holds. For the induction step, suppose that (7) holds for any $T \subsetneq R$. Take any $i \in R$, and let $T = R \setminus \{i\}$ and $Y = X \setminus \{i\}$. Let $C \in \mathcal{C}(g - Y)$ be the component with $i \in C$. Since R does not contain an essential separator of $g - X$ so $\mathcal{C}(g - X)$ and $\mathcal{C}(g - Y)$ differ at component C only: either $C \setminus \{i\} \in \mathcal{C}(g - X)$ or $C \setminus \{i\} = \emptyset$. Hence

$$\Phi(g - X) = \Phi(g - Y) - (f(|C|) - f(|C| - 1)). \tag{8}$$

Now let $C' \in \mathcal{C}(g - T)$ be the component with $i \in C'$. Applying attack $\{i\}$ to $g - T$ replaces C' with components C'_1, \dots, C'_m such that $\bigcup_{i=1}^m C'_i = C' \setminus \{i\}$. Hence

$$\begin{aligned} \Phi(g - R) &= \Phi(g - T) - \left(f(|C'|) - \sum_{i=1}^m f(|C'_i|) \right) \\ &\leq \Phi(g - T) - (f(|C'|) - f(|C'| - 1)) \end{aligned} \tag{9}$$

(by the fact that f is strictly convex). By the induction hypothesis,

$$\begin{aligned} \Phi(g - E) - \Phi(g - Y) + (f(|C|) - f(|C| - 1)) \\ \leq \Phi(g) - \Phi(g - T) + (f(|C|) - f(|C| - 1)) \end{aligned}$$

and, by the fact that $C \subseteq C'$ and by convexity of f ,

$$\begin{aligned} \Phi(g - E) - \Phi(g - Y) - (f(|C|) - f(|C| - 1)) \\ \leq \Phi(g) - \Phi(g - T) + (f(|C'|) - f(|C'| - 1)). \end{aligned}$$

Thus, by (8) and (9),

$$\Phi(g - E) - \Phi(g - X) \leq \Phi(g) - \Phi(g - R).$$

This shows the induction step. Hence we have shown (7).

Now, let $\Delta \subseteq N$ be a defense chosen in the first stage and suppose that X is a best response to Δ . Whereas X is a better response to Δ than R , so

$$-\Phi(g - X) - c_A|X| \geq -\Phi(g - R) - c_A|R|$$

and, consequently,

$$\Phi(g - R) \geq \Phi(g - X) + c_A(|X| - |R|) = \Phi(g - X) + c_A|E|.$$

From (7), we have

$$\Phi(g - X) \geq \Phi(g - E) + \Phi(g - R) - \Phi(g).$$

Putting the last two inequalities together, we arrive at

$$\Phi(g - R) \geq \Phi(g - E) + \Phi(g - R) - \Phi(g) + c_A|E|.$$

Simplifying this yields

$$-\Phi(g - E) - c_A|E| \geq -\Phi(g).$$

In other words, $E \in \mathcal{E}(g, c_A)$. □

The proof of part (ii) of Proposition 2 now follows from the lemmata above and the arguments in the main text. We turn next to proving part (i) of Proposition 2.

To simplify some parts of the argument, we will make a tie-breaking assumption on the behavior of the adversary. It says that if two strategies yield equal payoffs to the adversary, then he will choose the strategy that yields a lower payoff to the defender.

ASSUMPTION 2. *Given a network g and defense Δ , if two strategies $X \subseteq N$ and $X' \subseteq N$ yield the same payoff to the adversary, then he chooses the strategy that results in a residual network of lower value.*

The first step here is to state and prove the following lemma.

LEMMA 6. *Let $g \in \mathcal{G}(N)$ be a connected network over N , and let c_D and c_A be the costs of defense and attack, respectively. Suppose that $\Delta \subseteq N$ is an equilibrium defense and $X \subseteq N$ is a best response to it. Suppose that there exists $i \in \Delta$ such that $\mathcal{D}(\Delta, \mathcal{E}(g, c_A)) = \mathcal{D}(\Delta \setminus \{i\}, \mathcal{E}(g, c_A))$. Let $X' \subseteq N$ be a best response to $\Delta' = \Delta \setminus \{i\}$.*

Then there exists a component $C \in \mathcal{C}(g - X)$ such that $C \subseteq \Delta$ and either $C = \{i\}$ or $C \setminus \{i\} \in \mathcal{C}(g - X')$. Moreover,

$$\Pi^D(\Delta, X; g) = \Pi^D(\Delta', X'; g) + f(|C|) - f(|C| - 1) - c_D \tag{10}$$

and

$$c_A \leq f(|C|) - f(|C| - 1). \tag{11}$$

PROOF. Let $\Delta \subseteq N$ be a defense, $i \in \Delta$ and $\Delta' = \Delta \setminus \{i\}$. Let X be a best response to Δ and let X' be a best response to Δ' .

Since X is a best response to Δ , so $X \cap \Delta = \emptyset$ and $\Phi(g - (X \setminus \Delta)) = \Phi(g - X)$, and analogously with X' and Δ' . We prove the lemma in the seven steps below.

(i) We have $\Phi(g - X) > \Phi(g - X')$. Since Δ is an equilibrium strategy of the defender, so $\Pi^D(\Delta, X; g) \geq \Pi^D(\Delta', X'; g)$, that is, $\Phi(g - X) - c_D|\Delta| \geq \Phi(g - X') - c_D(|\Delta| - 1)$. Hence $\Phi(g - X) > \Phi(g - X')$.

(ii) We have $i \in X'$. Assume, to the contrary, that $i \notin X'$. Then $X' \cap \Delta = X' \cap \Delta' = \emptyset$. Similarly, since $X \cap \Delta = \emptyset$, so $X \cap \Delta' = \emptyset$. Hence $\Pi^A(\Delta', X'; g) = \Pi^A(\Delta, X'; g)$ and $\Pi^A(\Delta', X; g) = \Pi^A(\Delta, X; g)$. By the fact that $\Pi^A(\Delta', X'; g) \geq \Pi^A(\Delta', X; g)$, as X' is a best response to Δ' , this yields $\Pi^A(\Delta, X'; g) \geq \Pi^A(\Delta, X; g)$. Additionally, by point (i), $\Phi(g - X) > \Phi(g - X')$, so X' results in a residual network of lower value than in the case of X . Hence, by the tie-breaking Assumption 2, X' is an equilibrium response to Δ , a contradiction. Thus it must be that $i \in X'$.

(iii) Let $Y \subseteq N$ be an attack with $i \in Y$ and such that $Y \cap \Delta' = \emptyset$. Then either $\{i\} \cup \partial_g(i) \subseteq Y$ or there exists exactly one $C \in \mathcal{C}(g - Y)$ such that $i \in \partial_g(C)$.

Take any decomposition $E \cup R$ of Y , as described in Lemma 1. It cannot be that $i \in E$, as otherwise we would have $E \in \mathcal{D}(\Delta, \mathcal{E}(g, c_A))$, while $E \notin \mathcal{D}(\Delta', \mathcal{E}(g, c_A))$, as $Y \cap \Delta' = \emptyset$, and we would have a contradiction with the assumption that $\mathcal{D}(\Delta, \mathcal{E}(g, c_A)) = \mathcal{D}(\Delta', \mathcal{E}(g, c_A))$. Hence $i \in R$ and there exists a component $\tilde{C} \in \mathcal{C}(g - E)$ such that $i \in \tilde{C}$. Let $C = \tilde{C} \setminus R$ be what remains of \tilde{C} after the remainder R of Y is applied to $g - E$. By the definition of the remainder, R does not contain a separator for $g - E$. Therefore either $C = \emptyset$ (i.e., it is completely removed by R) or $C \in \mathcal{C}(g - Y)$ (i.e., it is a component in $g - Y$). Suppose that $C = \emptyset$, that is, $C \subseteq R$. Then $\partial_{g-E}(i) \subseteq R$ and $\partial_g(i) \subseteq E \cup R = Y$. Since $i \in Y$, so $\{i\} \cup \partial_g(i) \subseteq Y$. Suppose now that C is a component in $\mathcal{C}(g - Y)$. We will show that $i \in \partial_{g-E}(C)$. Assume the opposite. Then $\partial_{g-E}(C)$ must be a separator in $g - E$, as it separates C from a component containing i . But then $\partial_{g-E}(C)$ contains an essential separator for $g - E$. Since $\partial_{g-E}(C) \subseteq R$, this contradicts the assumption that R is a remainder and does not contain any essential separators of $g - E$. Hence it must be that $i \in \partial_{g-E}(C)$ and, consequently, $i \in \partial_g(C)$.

(iv) For all $C' \in \mathcal{C}(g - X')$ with $i \in \partial_g(C')$, $C' \subseteq \Delta$. Assume the opposite. Then there exists $C' \in \mathcal{C}(g - X')$ with $i \in \partial_g(C')$ (and consequently $i \notin C'$) such that $i' \in C' \setminus \Delta$. Consider a strategy $X'' = (X' \setminus \{i\}) \cup \{i'\}$. Since $X' \cap \Delta' = \emptyset$ and $i' \notin \Delta$, so $X'' \cap \Delta' = \emptyset$. Notice that $\Phi(g - X'') \leq \Phi(g - X')$, as both the residual networks agree at all the components apart from what remains of $C' \cup \{i\}$ after i' is removed (at the least it is one component of the same size as C'). Since $|X'| = |X''|$ so $\Pi^A(\Delta', X''; g) \geq \Pi^A(\Delta', X'; g)$ and so X'' is a best response to Δ' . But then we get a contradiction with point (ii), as $i \notin X''$. Hence it must be that $C' \subseteq \Delta$.

(v) There exists $C' \in \mathcal{C}(g - X') \cup \{\emptyset\}$ such that $C = C' \cup \{i\} \in \mathcal{C}(g - X)$ and $C \subseteq \Delta$. Let $C' = \emptyset$ if $\{i\} \cup \partial_g(i) \subseteq X'$ or let C' be the unique $C' \in \mathcal{C}(g - X')$ with $i \in \partial_g(C')$, otherwise. By point (iii) such C' exists. By point (iv) and by the fact that $i \in \Delta$, $C \subseteq \Delta$. Thus there exists a component $C'' \in \mathcal{C}(g - X)$ such that $C \subseteq C''$. Suppose that $C \subsetneq C''$. We will show that in this case $X \cup \{i\}$ is a better response to Δ' than X' , a contradiction.

Notice that since $X \cap \Delta = \emptyset$ and $\Delta' = \Delta \setminus \{i\}$ so $(X \cup \{i\}) \cap \Delta' = \emptyset$. By point (iii) either $\{i\} \cup \partial_g(i) \subseteq X \cup \{i\}$ or there exists exactly one component $C''' \in \mathcal{C}(g - (X \cup \{i\}))$ such that $i \in \partial_g(C''')$. Hence $C'' = C''' \cup \{i\}$ and C'' must be unique in $\mathcal{C}(g - X)$ with $i \in \partial_g(C'')$. The residual network $g - (X \cup \{i\})$ differs from $g - X$ at one component only: instead of C'' it has $C'' \setminus \{i\}$. Thus the value of residual network $g - (X \cup \{i\})$ is

$$\Phi(g - (X \cup \{i\})) = \Phi(g - X) - (f(|C''|) - f(|C''| - 1)). \tag{12}$$

Similarly, since either $C' = \emptyset$ or $i \in \partial_g(C')$, so the residual network when using $X' \setminus \{i\}$ against Δ' , $g - (X' \setminus \{i\})$, differs from $g - X'$ by one component: it has C instead of C' . Additionally, since $\Delta = \Delta' \cup \{i\}$ and $X' \cap \Delta' = \emptyset$ so $X' \setminus \{i\} = X' \setminus \Delta$. Thus the value of the residual network $g - (X' \setminus \{i\})$ can be written as

$$\Phi(g - (X' \setminus \{i\})) = \Phi(g - X') + f(|C|) - f(|C| - 1). \tag{13}$$

Since X is a best response to Δ , it is not worse than $X' \setminus \{i\}$. Hence

$$-\Phi(g - X) - c_A|X| \geq -\Phi(g - (X' \setminus \{i\})) - c_A(|X'| - 1).$$

This, together with (13), implies

$$\Phi(g - X) \leq \Phi(g - X') + f(|C|) - f(|C| - 1) - c_A(|X| - |X'| + 1). \tag{14}$$

Similarly, since X' is a best response to Δ' , it is not worse than $X \cup \{i\}$. Hence

$$-\Phi(g - X') - c_A|X'| \geq -\Phi(g - (X \cup \{i\})) - c_A(|X| + 1).$$

This, together with (12), implies

$$\Phi(g - X) \geq \Phi(g - X') + (f(|C''|) - f(|C''| - 1)) - c_A(|X| - |X'| + 1). \tag{15}$$

From (14) and (15) we get

$$\begin{aligned} f(|C''|) - f(|C''| - 1) - (f(|C|) - f(|C| - 1)) \\ \leq c_A(|X| + 1) - c_A|X| - (c_A|X'| - c_A(|X'| - 1)) = 0. \end{aligned}$$

If $C \subsetneq C''$, then $|C| < |C''|$, and, by strict convexity of f , the left-hand side is greater than 0, a contradiction. Thus it must be that $C'' = C$.

(vi) We have $\Pi^D(\Delta, X; g) = \Pi^D(\Delta', X'; g) + f(|C|) - f(|C| - 1) - c_D$. Since X is a best response to Δ , it is not worse than $X' \setminus \{i\}$. Hence

$$-\Phi(g - X) - c_A|X| \geq -\Phi(g - (X' \setminus \{i\})) - c_A(|X'| - 1).$$

Adding $f(|C|) - f(|C| - 1)$ to both sides we get

$$\begin{aligned} -(\Phi(g - X) - (f(|C|) - f(|C| - 1))) - c_A|X| \\ \geq -(\Phi(g - (X' \setminus \{i\})) - (f(|C|) - f(|C| - 1))) - c_A(|X'| - 1). \end{aligned} \tag{16}$$

As we observed in the proof of point (v) ((12) and (13) and the fact that $C'' = C$),

$$\Phi(g - (X \cup \{i\})) = \Phi(g - X) - (f(|C|) - f(|C| - 1)) \tag{17}$$

$$\Phi(g - X') = \Phi(g - (X' \setminus \{i\})) - (f(|C|) - f(|C| - 1)). \tag{18}$$

Hence, from (16), we get

$$-\Phi(g - (X \cup \{i\})) - c_A(|X| + 1) \geq -\Phi(g - X') - c_A|X'|.$$

Alternatively, since X' is a best response to Δ' , so

$$-\Phi(g - (X \cup \{i\})) - c_A(|X| + 1) \leq -\Phi(g - X') - c_A|X'|.$$

Combining these two inequalities we get

$$-\Phi(g - (X \cup \{i\})) - c_A(|X| + 1) = -\Phi(g - X') - c_A|X'|. \tag{19}$$

Since X' is the equilibrium response to Δ' , by tie-breaking [Assumption 2](#),

$$\Phi(g - X') \leq \Phi(g - (X \cup \{i\})).$$

Additionally this, together with (17) and (18), implies

$$\Phi(g - (X' \setminus \{i\})) \leq \Phi(g - X). \tag{20}$$

From (19), (17), and (18) we get

$$-\Phi(g - X) - c_A|X| = -\Phi(g - (X' \setminus \{i\})) - c_A(|X'| - 1).$$

Again, since X is the equilibrium response to Δ , by tie-breaking [Assumption 2](#),

$$\Phi(g - X) \leq \Phi(g - (X' \setminus \{i\})),$$

and, by (17), (18), and (20),

$$\begin{aligned} \Phi(g - X) &= \Phi(g - (X' \setminus \{i\})), \\ \Phi(g - (X \cup \{i\})) &= \Phi(g - X'). \end{aligned}$$

Thus both X and $X' \setminus \{i\}$ are best responses to Δ and both X' and $X \cup \{i\}$ are best responses to Δ' . This, together with (18), implies

$$\Pi^D(\Delta, X; g) = \Pi^D(\Delta', X'; g) + f(|C|) - f(|C| - 1) - c_D.$$

(vii) We have $c_A \leq f(|C|) - f(|C| - 1)$. Since X' is a better response to Δ' than $X' \setminus \{i\}$, so

$$-\Phi(g - X') - c_A|X'| \geq -\Phi(g - (X' \setminus \{i\})) - c_A(|X'| - 1)$$

and, consequently,

$$c_A \leq \Phi(g - (X' \setminus \{i\})) - \Phi(g - X').$$

By (17),

$$c_A \leq f(|C|) - f(|C| - 1). \quad \square$$

PROOF OF PART (I) OF PROPOSITION 2. Characterization of the optimal strategies of the adversary follows directly from Lemma 1. Thus in what follows we concentrate on the equilibrium defense.

Let Δ be an equilibrium defense. We will show first that if $\Delta \subsetneq N$, then Δ must be a minimal transversal of $\mathcal{D}(\Delta, \mathcal{E}(g, c_A))$.

Assume the opposite. Then there exists $i \in \Delta$ such that $\mathcal{D}(\Delta \setminus \{i\}, \mathcal{E}(g, c_A)) = \mathcal{D}(\Delta, \mathcal{E}(g, c_A))$. Let X be the equilibrium response to Δ and let X' be the equilibrium response to $\Delta' = \Delta \setminus \{i\}$. Clearly $X \cap \Delta = \emptyset$ and $X' \cap \Delta' = \emptyset$.

Recall that $\mathcal{C}(g - X')$ is the set of components in the residual network when the strategies Δ' and X' are used by the players, and $\mathcal{C}(g - X)$ is the set of components in the residual network when Δ and X are used. By the assumption that $\Delta \subsetneq N$, both these sets are nonempty. We will show that either Δ' or Δ'' (described below) is a better strategy for the defender than Δ , which will contradict the assumption that Δ is an equilibrium strategy.

Let $C \in \mathcal{C}(g - X)$ be a component such that $C \subseteq \Delta$ and either $C = \{i\}$ or $C \setminus \{i\} \in \mathcal{C}(g - X')$. By Lemma 6 such C exists.

Since for all $j \in \partial_g(C)$, $\mathcal{D}(\Delta, \mathcal{E}(g, c_A)) \subsetneq \mathcal{D}(\Delta \cup \{j\}, \mathcal{E}(g, c_A))$, any such j belongs to an essential separator not covered by Δ . Take any $j \in \partial_g(C)$ and let $\{C_1, \dots, C_m\} \subseteq \mathcal{C}(g - X)$ be all the components in $g - X$ such that $j \in \partial_g(C_l)$ for all $l \in \{1, \dots, m\}$ (assume, without loss of generality, that $C_1 = C$; notice that in particular it may be that $m = 1$ and the argument below works for that case as well). Consider defenses $\Delta' = \Delta \setminus \{i\}$ and $\Delta'' = \Delta \cup \{j\} \cup \bigcup_{l=2}^m C_l$. We will show that either Δ' or Δ'' is a better strategy for the defender than Δ .

Let X'' be the equilibrium response of the adversary to Δ'' and let $C'' = \{j\} \cup \bigcup_{l=1}^m C_l$. We show first that $C'' \in \mathcal{C}(g - X'')$. Since Δ'' protects C'' , there is component $C''' \in \mathcal{C}(g - X'')$ such that $C'' \subseteq C'''$. Suppose that $C'' \subsetneq C'''$. Then there exists $v \in C'''$ such that $v \notin C''$. We will show that $v \notin \Delta''$. If $v \in \partial_g(C_l)$ for some $l \in \{1, \dots, m\}$, then it cannot be that $v \in \Delta$ (because these components are separated by X used as an equilibrium response to Δ). Thus the only possibility is that $v \in \partial_g(\{j\})$. But then v would be one of the components C_l created by applying X to g and, consequently, it would be $v \in C''$, a contradiction with the assumption that $v \notin C''$. Since $v \notin \Delta$ and $v \notin C''$, then $v \notin \Delta''$. Now consider a response $X'' \cup \{v\}$ to Δ'' . At the very least it removes a node from component C''' (it may additionally disconnect the component). Hence $\Phi(g - (X'' \cup \{v\})) \leq \Phi(g - X'') - f(|C'''|) + f(|C'''| - 1)$. Alternatively, by Lemma 6, (11), $c_A \leq f(|C|) - f(|C| - 1) < f(|C'''|) - f(|C'''| - 1)$ (by convexity of f and $|C'''| \geq |C| + 1$). Thus it follows that

$$-\Phi(g - (X'' \cup \{v\})) - c_A(|X''| + 1) > -\Phi(g - X'') - c_A|X''|,$$

which contradicts the assumption that X'' is a best response to Δ'' . Therefore it must be $C''' = C''$.

As we have shown above, $C'' = \{j\} \cup \bigcup_{l=1}^m C_l \in \mathcal{C}(g - X'')$. After attack $X'' \cup \{j\}$ is applied to g , component C'' is replaced with components C_1, \dots, C_m . Hence

$$\Phi(g - X'') = \Phi(g - (X'' \cup \{j\})) + f\left(1 + \sum_{l=1}^m |C_l|\right) - \left(\sum_{l=1}^m f(|C_l|)\right). \tag{21}$$

Alternatively, since C is a component in $g - X$, every node in $\partial_g(C)$ is removed by X . Thus, when nodes in $\Delta \cup \{j\} \cup \bigcup_{l=2}^m C_l$ are defended, the residual network $g - (X \setminus \{j\})$ differs from $g - X$ by having component C'' instead of components C_1, \dots, C_m . Hence

$$\Phi(g - (X \setminus \{j\})) = \Phi(g - X) + f\left(1 + \sum_{l=1}^m |C_l|\right) - \left(\sum_{l=1}^m f(|C_l|)\right). \tag{22}$$

Since X'' is a better response to Δ'' than $X \setminus \{j\}$,

$$-\Phi(g - X'') - c_A |X''| \geq -\Phi(g - (X \setminus \{j\})) - c_A (|X| - 1) \tag{23}$$

and $\Phi(g - X'') \leq \Phi(g - (X \setminus \{j\}))$ in the case of equality (notice that $(X \setminus \{j\}) \cap \Delta'' = \emptyset$ as $X \cap C_l = \emptyset$ for all $l \in \{1, \dots, m\}$ and $X \cap \Delta = \emptyset$).

Equations (21), (22), and (23) imply

$$-\Phi(g - (X'' \cup \{j\})) - c_A |X''| \geq -\Phi(g - X) - c_A (|X| - 1).$$

Subtracting c_A from both sides we get

$$-\Phi(g - (X'' \cup \{j\})) - c_A (|X''| + 1) \geq -\Phi(g - X) - c_A |X|. \tag{24}$$

Alternatively, since X is a best response to Δ than is $X'' \cup \{j\}$, we have

$$-\Phi(g - X) - c_A |X| \geq -\Phi(g - (X'' \cup \{j\})) - c_A (|X''| + 1) \tag{25}$$

and $\Phi(g - X) \leq \Phi(g - (X'' \cup \{j\}))$, in the case of equality.

By (24) and (25), $X'' \cup \{j\}$ is a best response to Δ as well, and since X is an equilibrium response to Δ , it must be that $\Phi(g - X) \leq \Phi(g - (X'' \cup \{j\}))$. Combining this with (21) we get

$$\Phi(g - X'') \geq \Phi(g - X) + f\left(1 + \sum_{l=1}^m |C_l|\right) - \left(\sum_{l=1}^m f(|C_l|)\right). \tag{26}$$

From (10) and (26) it follows that

$$\begin{aligned} \Pi^D(\Delta, X; g) &= \Pi^D(\Delta', X'; g) + f(|C|) - f(|C| - 1) - c_D \\ \Pi^D(\Delta'', X''; g) &\geq \Pi^D(\Delta, X; g) + f\left(1 + \sum_{l=1}^m |C_l|\right) - \left(\sum_{l=1}^m f(|C_l|)\right) - c_D. \end{aligned}$$

Since Δ is a better strategy than Δ' , then $f(|C|) - f(|C| - 1) \geq c_D$. Alternatively, since Δ is a better strategy than Δ'' , then $c_D \geq f(1 + \sum_{l=1}^m |C_l|) - (\sum_{l=1}^m f(|C_l|))$. Hence $f(|C|) - f(|C| - 1) \geq f(1 + \sum_{l=1}^m |C_l|) - (\sum_{l=1}^m f(|C_l|))$, which contradicts strict convexity of f .

Thus we have shown that $\Delta \subsetneq N$, and then Δ must be a minimal transversal of $\mathcal{D}(\Delta, \mathcal{E}(g, c_A))$. □

PROOF OF PROPOSITION 3. The nonmonotonicities have been established in the text. Here we establish monotonicity of the defender’s payoff in cost of attack and monotonicity of the adversary’s payoff in cost of defense.

We start with monotonicity of payoff to the defender in cost of attack. The argument here is straightforward in the generic case, where equilibrium payoffs are unique: suppose (Δ^*, X^*) is an equilibrium with network g and costs (c_A, c_D) . Let $c'_A > c_A$. If the defender retains defense strategy Δ^* , it must be the case that the attack strategy will be weakly smaller under high cost c'_A . This in turn implies that the defender’s payoff must be weakly larger if he maintains the original strategy Δ^* . So, in equilibrium under (c'_A, c_D) , he must also do better. However, the monotonicity holds for any values of the parameters. The problem here is the nonuniqueness of equilibrium payoffs. However, this is not a concern, because if this was the case, the more costly attacks would cease being equally good for the adversary as the less costly ones. The precise argument is as follows. Let c_A and c'_A be the costs of attack such that $c'_A > c_A$. Let (Δ^*, X^*) be an equilibrium under c_A and let (Δ^{**}, X^{**}) be an equilibrium under c'_A . Since $X^*(\Delta^*)$ is a best response to Δ^* under c_A , it is not worse than $X^{**}(\Delta^*)$; hence

$$-\Phi(g - X^*(\Delta^*)) - c_A|X^*(\Delta^*)| \geq -\Phi(g - X^{**}(\Delta^*)) - c_A|X^{**}(\Delta^*)|,$$

which yields

$$\Phi(g - X^*(\Delta^*)) - \Phi(g - X^{**}(\Delta^*)) \leq c_A(|X^{**}(\Delta^*)| - |X^*(\Delta^*)|). \tag{27}$$

Similarly, since $X^{**}(\Delta^*)$ is a best response to Δ^* under c'_A , it is not worse than $X^*(\Delta^*)$. This yields

$$\Phi(g - X^*(\Delta^*)) - \Phi(g - X^{**}(\Delta^*)) \geq c'_A(|X^{**}(\Delta^*)| - |X^*(\Delta^*)|). \tag{28}$$

Equations (27) and (28) imply $c_A(|X^{**}(\Delta^*)| - |X^*(\Delta^*)|) \geq c'_A(|X^{**}(\Delta^*)| - |X^*(\Delta^*)|)$. By $c'_A > c_A$ it follows that

$$|X^{**}(\Delta^*)| \leq |X^*(\Delta^*)|. \tag{29}$$

Now assume to the contrary that

$$\Pi^D(\Delta^*, X^*(\Delta^*); g, c_D) > \Pi^D(\Delta^{**}, X^{**}(\Delta^{**}); g, c_D).$$

Since Δ^{**} is an equilibrium defense under c'_A ,

$$\Pi^D(\Delta^{**}, X^{**}(\Delta^{**}); g, c_D) \geq \Pi^D(\Delta^*, X^{**}(\Delta^*); g, c_D).$$

The two equations above imply

$$\Pi^D(\Delta^*, X^*(\Delta^*); g, c_D) > \Pi^D(\Delta^*, X^{**}(\Delta^*); g, c_D),$$

that is,

$$\Phi(g - X^*(\Delta^*)) - c_D|\Delta^*| > \Phi(g - X^{**}(\Delta^*)) - c_D|\Delta^*|$$

and, consequently,

$$\Phi(g - X^*(\Delta^*)) - \Phi(g - X^{**}(\Delta^*)) > 0. \tag{30}$$

Equations (27) and (30) imply $c_A(|X^{**}(\Delta^*)| - |X^*(\Delta^*)|) > 0$. By $c_A > 0$, it follows that $|X^{**}(\Delta^*)| > |X^*(\Delta^*)|$, a contradiction with (29). Thus it must be that $\Pi^D(\Delta^*, X^*(\Delta^*); g, c_D) \leq \Pi^D(\Delta^{**}, X^{**}(\Delta^{**}); g, c_D)$. Notice that this argument holds for any parameters of the model, not only in the generic case.

We now turn to the monotonicity of payoff to the adversary in cost of defense. Let c_D and c'_D be the costs of defense such that $c'_D > c_D$. Let (Δ^*, X^*) be an equilibrium under c_D and let (Δ^{**}, X^{**}) be an equilibrium under c'_D . Since $X^*(\Delta^*)$ is a best response to Δ^* and $X^{**}(\Delta^*)$ is a best response to Δ^* in the adversary’s subgame,

$$-\Phi(g - X^*(\Delta^*)) - |X^*(\Delta^*)|c_A = -\Phi(g - X^{**}(\Delta^*)) - |X^{**}(\Delta^*)|c_A.$$

Thus another equilibrium under c_D is (Δ^*, X') , where X' equals X^* at all defense profiles but Δ^* , where it is equal to Δ^{**} . By Lemma 3, generically, $\Phi(g - X^*(\Delta^*)) = \Phi(g - X^{**}(\Delta^*))$. By analogous arguments, $\Phi(g - X^{**}(\Delta^{**})) = \Phi(g - X^*(\Delta^{**}))$. Since Δ^* is an equilibrium defense under c_D and Δ^{**} is an equilibrium defense under c'_D ,

$$\begin{aligned} \Phi(g - X^*(\Delta^*)) - |\Delta^*|c_D &\geq \Phi(g - X^*(\Delta^{**})) - |\Delta^{**}|c_D \\ \Phi(g - X^{**}(\Delta^{**})) - |\Delta^{**}|c'_D &\geq \Phi(g - X^{**}(\Delta^*)) - |\Delta^*|c'_D, \end{aligned}$$

which can be rewritten as

$$\begin{aligned} \Phi(g - X^*(\Delta^{**})) - \Phi(g - X^*(\Delta^*)) &\leq (|\Delta^{**}| - |\Delta^*|)c_D \\ \Phi(g - X^{**}(\Delta^{**})) - \Phi(g - X^{**}(\Delta^*)) &\geq (|\Delta^{**}| - |\Delta^*|)c'_D. \end{aligned}$$

Since $c'_D > c_D$, these inequalities imply

$$\Phi(g - X^{**}(\Delta^{**})) - \Phi(g - X^{**}(\Delta^*)) > \Phi(g - X^*(\Delta^{**})) - \Phi(g - X^*(\Delta^*)).$$

This, combined with $\Phi(g - X^*(\Delta^*)) = \Phi(g - X^{**}(\Delta^*))$ and $\Phi(g - X^{**}(\Delta^{**})) = \Phi(g - X^*(\Delta^{**}))$, leads to contradiction. Hence it must be that the payoff to the adversary increases when the cost of defense increases. Notice that this argument holds for generic values of the parameters of the model. There are nongeneric examples where the payoff to the adversary decreases when the cost of defense increases. \square

Before proving Proposition 4, we need the following auxiliary lemma, stating a useful property of a convex function.

LEMMA 7. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a strictly convex and differentiable function. Then function*

$$g(x, y) = \frac{yf(x) - xf(y)}{x - y}$$

is strictly increasing in both arguments as long as $x > y$.

PROOF. To show the result we compute partial derivatives of h :

$$g_x(x, y) = \left(\frac{y}{x-y}\right) \left(f'(x) - \frac{f(x) - f(y)}{x-y}\right)$$

$$g_y(x, y) = \left(\frac{x}{x-y}\right) \left(\frac{f(x) - f(y)}{x-y} - f'(y)\right).$$

By strict convexity of f , $f'(y) < (f(x) - f(y))/(x - y) < f'(x)$ as long as $x > y$; hence $g_x, g_y > 0$ and g is strictly increasing in x and in y . This completes the proof. \square

Now we are ready to prove Proposition 4.

PROOF OF PROPOSITION 4. Point (i) follows directly and we omit the proof. For point (ii) observe, from Proposition 2, that with $c_A \in (\Delta f(n - 1), f(n) - (n - 1)f(1))$ and $g \in \mathcal{E}(g, c_A) = 0$, the optimal attack targets no nodes. So the optimal defense also consists of defending no nodes. Thus the costs of conflict are 0.

For point (iii), assume that $c_A \in (\Delta f(m - 1), \Delta f(m))$ with $m \in \{1, \dots, n - 1\}$. With such a cost of attack, on any connected network, the adversary best responds to any incomplete defense by removing at least one node. Therefore the lower bound for the costs of conflict are $\min(c_A, nc_D)$ in this case.

Part 1. Suppose that $c_D > c(n, m)$. We show first that in every equilibrium on h_n^m the defender chooses the empty defense and the adversary responds to it with attack $\{1\}$ (the separator of h_n^m). By Proposition 2, an equilibrium defense must be either empty, or complete, or equal to $\{1\}$. Moreover, the best response of the adversary to the empty defense either contains $\{1\}$, in which case the reducing attack part of it must be empty (because components of $h_n^m - \{1\}$ have sizes at most m), or does not contain $\{1\}$, in which case it must be a reducing attack leaving a residual network consisting of a single component of size m . It is easy to check that the former is the best response to the empty defense and the latter is the best response to defense $\{1\}$. Hence empty defense is better than $\{1\}$. The payoff to the defender from using the empty defense is

$$\Pi^D(\emptyset, \{1\}; h_n^m, c_D) = \Phi(h_n^m - \{1\}) = \left\lfloor \frac{n-1}{m} \right\rfloor f(m) + f((n-1) \bmod m).$$

With cost of defense $c_D > c(m, n)$, the payoff to the defender from the complete defense,

$$\Pi^D(N, \emptyset; h_n^m, c_D) = f(n) - nc_D,$$

is lower than the payoff from the empty defense. Hence on the equilibrium path the defender chooses \emptyset and the adversary responds with $\{1\}$.

Second, we show that for the ranges of costs in question, $nc_D > c_A$. Since $c_D > c(n, m)$,

$$nc_D > f(n) - \left\lfloor \frac{n-1}{m} \right\rfloor f(m) - f((n-1) \bmod m).$$

The right-hand side of this inequality can be rewritten as

$$f(n) - \frac{n-1 - (n-1) \bmod m}{m} f(m) - f((n-1) \bmod m).$$

Since f is strictly convex and $(n - 1) \bmod m < m$, then $((n - 1) \bmod m)f(m) > mf((n - 1) \bmod n)$. Therefore,

$$c_D > f(n) - \left(\frac{n-1}{m}\right)f(m).$$

The right-hand side can be rewritten as

$$\begin{aligned} f(n) - \left(\frac{n-1}{m}\right)f(m) &= \sum_{j=m}^{n-1} \Delta f(j) - (n-m-1)\left(\frac{f(m)}{m}\right) \\ &= \Delta f(m) + \sum_{j=m+1}^{n-1} \left(\Delta f(j) - \frac{f(m)}{m}\right). \end{aligned}$$

By convexity of f , for all $j > m$, $\Delta f(j) > f(m)/m$. Thus $nc_D > \Delta f(m)$ and, since $c_A \in (\Delta f(m - 1), \Delta f(m))$, $nc_D > c_A$. Hence the minimal costs of conflict are c_A .

Part 2. Suppose that $c_D < c(n, m)$. We will show that with such a cost of defense, in any equilibrium on a connected network the defender chooses the complete defense. Notice that with $c_D < c(n, m)$, on any connected network g , any defense Δ of size $|\Delta| \leq m$ is worse for the defender than the complete defense. This is because the residual network after the adversary best responding to Δ consists of components of sizes at most m and the upper bound on the value of such residual networks is $\lfloor (n - 1)/m \rfloor f(m) + f((n - 1) \bmod m)$ (this upper bound is attained by h_n^m). With $c_D < c(n, m)$ the defender prefers complete defense to Δ .

Consider defense Δ of size $d = |\Delta|$ such that $m < d < n$. Let X be a best response to Δ . The payoff to the defender from Δ and X is

$$\begin{aligned} \Pi^D(\Delta, X; g, c_D) &= \Phi(g - X) - dc_D \\ &\leq f(d) + \left\lfloor \frac{n-d-1}{m} \right\rfloor + f((n-d-1) \bmod m) - dc_D. \end{aligned}$$

The upper bound on the value of the residual network above comes from the following observation. With $c_A \in (\Delta f(m - 1), \Delta f(m))$, in any best response the adversary removes unprotected nodes from any component of size greater than m . Therefore in the best case the adversary removes one node and the only component of size greater than m in the residual network is a fully protected component of size d (by convexity of f it is better to have one fully protected component of size d than several fully protected and smaller ones summing up to d). Thus if

$$c_D < \frac{f(n) - f(d) - \left\lfloor \frac{n-d-1}{m} \right\rfloor - f((n-d-1) \bmod m)}{n-d},$$

then the complete defense is better to Δ for the defender. We will show that $c(n, m)$ is lower than the right-hand side of the inequality above, which will imply that for the costs of defense under consideration, the complete defense is better for the defender.

The inequality

$$c(n, m) = \frac{f(n) - \lfloor \frac{n-1}{m} \rfloor - f((n-1) \bmod m)}{n} < \frac{f(n) - f(d) - \lfloor \frac{n-d-1}{m} \rfloor - f((n-d-1) \bmod m)}{n-d}$$

can be rewritten as

$$df(n) - nf(d) - \left(\frac{n-d}{m}\right)(r_1f(m) - mf(r_1)) + \left(\frac{n}{m}\right)(r_2f(m) - mf(r_2)) > 0,$$

where $r_1 = (n-1) \bmod m$ and $r_2 = (n-d-1) \bmod m$.⁹ Since $r_2 < m$ and f is convex, then $r_2f(m) - mf(r_2) > 0$, and to show that the inequality above holds it suffices to show that

$$df(n) - nf(d) - \left(\frac{n-d}{m}\right)(r_1f(m) - mf(r_1)) > 0. \tag{31}$$

Since $d < n$ and f is convex, $df(n) - nf(d) > 0$. Moreover, by Lemma 7,

$$\frac{df(n) - nf(d)}{n-d} > \frac{r_1f(m) - mf(r_1)}{m-r_1}$$

(as $n > d > m > r_1$). Hence

$$\frac{df(n) - nf(d)}{n-d} - \left(\frac{m-r_1}{m}\right)\left(\frac{r_1f(m) - mf(r_1)}{m-r_1}\right) > 0,$$

which implies (31), by multiplying both sides by $(n-d)$. Hence any equilibrium defense is complete and the costs of conflict are nc_D . □

Example of no conflict networks

Even if the marginals of f do not grow very rapidly, there may exist networks (other than complete network) that do not feature active conflict. Take $f(x) = x^2$, for example. Consider a family of core-periphery networks, $\{cp_k\}_{k \in \mathbb{N}}$. Given $k \in \mathbb{N}$, network cp_k has $2k$ nodes: a fully connected core of k nodes, and a periphery of k nodes. Each core node is connected to exactly one, unique, periphery node (cf. Figure 9).

When the cost of attack is high, $c_A > 4m - 1$, then it is easy to verify that in equilibrium the defender will either defend all the core nodes or use an empty defense. When the cost of attack is low, $c_A < 4m - 1$, then, again, there are two types of equilibrium defense: either no node is defended or all nodes are defended. It is easy to verify that three types of defense would be candidates for equilibrium defense here: empty defense, complete defense, and defense with all core nodes protected. To rule out the last one, suppose that $2(2m - k) - 1 \leq c_A < 2(2m - k) + 1$, where $1 \leq k \leq m - 1$.

⁹Recall that for integer x and y , $\lfloor x/y \rfloor = (x - x \bmod y)/y$.




FIGURE 9. Core-periphery networks cp_2 , cp_4 , and cp_6 .




FIGURE 10. Network that allows for active conflict (under $f(x) = x^2$).

With such a cost of attack, the adversary would remove a single node from a component of size greater than $2m - k$ and would not remove single nodes from a component of size $2m - k$ or less. Thus, when all the core nodes are protected, the adversary removes k spokes. When no node is protected, the adversary removes k core nodes. Since defending only the core nodes is better than defending all the nodes, $mc_D \geq (2m)^2 - (2m - k)^2$. Alternatively, protecting all the core nodes is better than no protection, so $mc_D \leq (2m - k)^2 - (2m - 2k)^2 - k$. The two inequalities imply $2k^2 - k \leq 0$, a contradiction for $k \geq 1$.

Notice, in the example above, that if each core node was connected to a higher number of periphery nodes, active conflict would be possible (as illustrated by Example 1). With more periphery nodes per core node (and with suitable costs of defense and attack), protecting the separators may create enough value for such a defense to be attractive. Increasing the value of the residual network requires defending all the nodes, which is too high an investment and too low a gain to be profitable. This illustrates one reason for the possibility of active conflict in the model: blocking all the individually rational essential separators may secure a high value of the residual network at a relatively low cost, while increasing the value further may require a much higher cost.

To get more insight into why active conflict is possible, despite the convexity of f and the linearity of costs, consider the network in Figure 10. Figure 11 illustrates function $\Phi^*(m; g, c_A)$ under different ranges of costs of attack. The dotted line is an upper convex hull of that function. The optimal size of defense is at a point of that hull adjacent to a line with slope c_D . In the case of low cost of attack, if the convex hull contains any points of $\Phi^*(m; g, c_A)$ for $0 < m < n$, then active conflict is possible for some suitable range of costs of defense. In the case of high cost of attack, active conflict is possible if the convex hull contains any points of $\Phi^*(m; g, c_A)$ for $0 < m < \tau(\mathcal{E}(g, c_A))$.

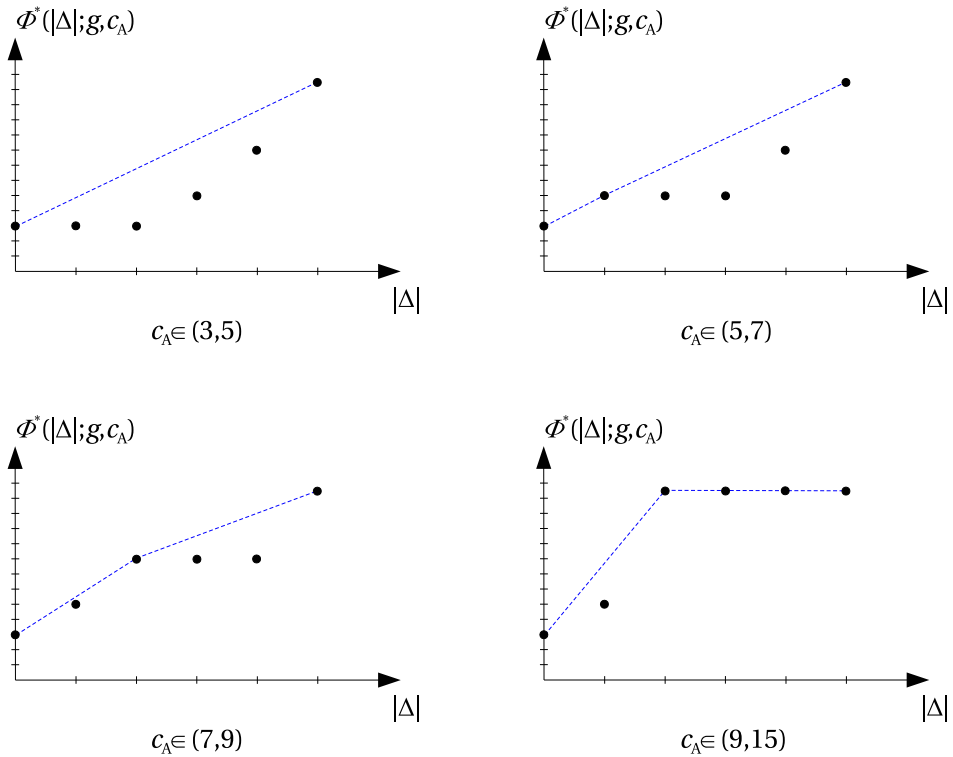


FIGURE 11. Optimal defenses of different sizes for network in Figure 10.

In Figure 11, low cost of attack is $c_A < 9$ and $c_A > 9$ is high cost of attack. Active conflict is possible for $c_A \in (5, 9)$. When $c_A \in (5, 7)$ and $c_D \in (3.75, 4)$, then the unique equilibrium defense is $\Delta^* = \{b\}$, and the best response to it in the adversary’s subgame is $X^*(\Delta^*) = \{a\}$. When $c_A \in (7, 9)$ and $c_D \in (5, 9)$, then the unique equilibrium defense is $\Delta^* = \{a, b\}$ and removing any unprotected node is a best response to it in the adversary’s subgame. When $c_A \in (9, 15)$, $\tau(\mathcal{E}(g, c_A)) = 2$ and there is no equilibrium outcome with active conflict.

PROOF OF PROPOSITION 5. For point (i), suppose that $c_D > f(n)/n$. We will show that in this case the equilibrium defense $\Delta = \emptyset$. Assume, to the contrary, that $\Delta \neq \emptyset$ and let X be the equilibrium response to Δ . Pick any $i \in \Delta$ and let $C(i)$ be the component of i in the residual network $g - X$. The payoff to i is

$$\Pi^i(\Delta, X; g, c_D) = \frac{f(|C(i)|)}{|C(i)|} - c_D.$$

By the fact that f is strictly increasing and strictly convex, $f(x)/x$ is increasing. Hence $\Pi^i(\Delta, X; g, c_D) \leq f(n)/n - c_D < 0$. Thus i is better off by not protecting, a contradiction to the assumption that Δ is an equilibrium defense. Hence it must be that $\Delta = \emptyset$.

For point (ii), suppose that $c_D < f(n)/n$. Assume that $c_A < f(n) - f(n - 1)$. We will show that $\Delta = N$ is an equilibrium defense. Assume otherwise. Then there exists $i \in \Delta$

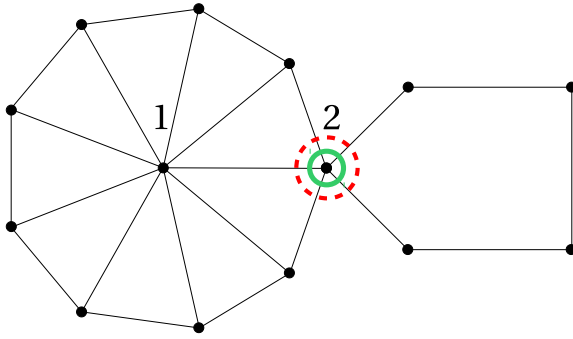


FIGURE 12. Separators and other centrality measures.

that is better off by deviating and choosing no protection. Since $c_A < f(n) - f(n - 1)$, the best response to $\Delta \setminus \{i\}$ is $X = \{i\}$, and so the deviating node gets removed, obtaining payoff 0 instead of $f(n)/n - c_D \geq 0$. Hence i is not better off by deviating and so $\Delta = N$ is an equilibrium defense. This proves point (a).

Assume that $c_A > f(n) - f(n - 1)$. Let Δ be minimal transversal of $\mathcal{E}(g, c_A)$. We will show that Δ is an equilibrium defense. By Lemma 1, the best response to Δ is the empty attack $X = \emptyset$. Assume, to the contrary, that Δ is not an equilibrium defense. Then there exists $i \in \Delta$ that is better off by choosing no protection instead of protection. Since Δ is a minimal transversal, it must be that there exists an essential separator $E \in \mathcal{E}(g, c_A)$ such that $\Delta \setminus \{i\} \cap E = \emptyset$. Moreover, any such separator contains i . Since any such separator is better than the empty attack, the adversary responds to $\Delta \setminus \{i\}$ with one of these separators, removing i . But then i gets payoff 0 instead of $f(n)/n - c_D \geq 0$. Hence it is not better off by deviating, a contradiction. Therefore Δ must be an equilibrium defense. This proves point (b).

Since the adversary’s subgame remains as in the centralized defense game, an equilibrium response X^* is as described in Proposition 2. \square

APPENDIX B: KEY PLAYERS AND CENTRALITY

Essential separators and their transversals determine the key nodes in our study of attack and defense. These key groups of nodes give rise to new notions of centrality distinct from other notions such as closeness, betweenness, or eigenvector centralities. To see how these notions are different, consider the network in Figure 12 (for simplicity the example is based on individual, rather than group, notions of centrality). Assume that the network value is based on function $f(x) = x^2$ and suppose that the cost of attack is $c_A \in (25, 89)$, so that the adversary attacks only the nodes that separate the network and so that removing node 2 is better than not attacking at all. Suppose also that $c_D \in (0, 89)$, so that defending node 2 constitutes an optimal defense as well. However, this node is less central than node 1 in the sense of degree, closeness, betweenness, eigenvector, Bonacich, and intercentrality measures.¹⁰ The numerical values for these centralities

¹⁰Following Ballester et al. (2006), we define for a parameter $\alpha \in \mathbb{R}$, $\mathbf{b}(g, \alpha) = \mathbf{M}(g, \alpha)\mathbf{1}$, where $\mathbf{M}(g, \alpha) = (\mathbf{I} - \alpha\mathbf{G})^{-1}$, \mathbf{I} is the identity matrix, and \mathbf{G} is the adjacency matrix of the network. We require α to be relatively

Centrality	Node 1	Node 2
Degree	9	5
Closeness	0.684	0.619
Betweenness	42.5	37.5
Eigenvector	0.5765	0.3036
Bonacich, high	532.2	281.18
Bonacich, medium	2.4311	1.8208
Bonacich, low	1.093	1.0519
Intercentrality, high	2940.9	2863.2
Intercentrality, medium	5.2438	3.1263
Intercentrality, low	1.1936	1.1061

TABLE 1. Centralities of nodes 1 and 2 in the network from Figure 12.




FIGURE 13. Separators and transversals in interlinked stars ($n = 12$).

are summarized in Table 1. For Bonacich centrality, we consider three values of the parameters: high ($\alpha = 0.237$), intermediate ($\alpha = 0.1$), and low ($\alpha = 0.01$).

APPENDIX C: SEPARATORS AND TRANSVERSALS IN FAMILIES OF NETWORKS

Interlinked stars

Interlinked stars are networks with two disjoint nonempty sets of nodes: the set of *centers* C and the set of *periphery nodes* P . The centers are fully connected, forming a clique. Each of the periphery nodes is connected to all the centers. Interlinked stars have one essential separator: the set of all the centers, $\mathcal{E}(g) = \{C\}$. All minimal transversals of $\mathcal{E}(g)$ are singleton sets consisting of one central node. The essential separator and a minimal transversal for an interlinked star are illustrated in Figure 13.

small so that $\mathbf{M}(g, \alpha)$ is well defined and nonnegative. The intercentrality measure we consider, also defined in that paper, is $c_i(g, \alpha) = b_i(g, \alpha)^2 / M_{ii}(g, \alpha)$. We define closeness as $cl_i(g) = (n - 1) / \sum_{j \neq i} d(i, j; g)$, where $d(i, j; g)$ is the length of the shortest path between i and j in g .




FIGURE 14. Separators and transversals in complete bipartite networks ($n = 12$).




FIGURE 15. Separators and transversals in trees ($n = 12$).

Complete bipartite networks

In a complete bipartite network the set of nodes, N , can be partitioned into two disjoint sets, N_1 and N_2 , $N_1 \cap N_2 = \emptyset$, such that the set of links is the set of all possible links connecting nodes from N_1 and nodes from N_2 . There are two essential separators in these networks, $\mathcal{E}(g) = \{N_1, N_2\}$. Every transversal consists of one node from N_1 and one node from N_2 . Minimal essential separators and transversals for complete bipartite networks are illustrated in Figure 14.

Trees

In any tree network, every nonempty set of internal nodes (nodes that are not leaves) constitutes a separator. Essential separators are sets of internal nodes such that no two of them are neighbors. Transversals of essential separators are subsets of internal nodes. In particular, there is a unique transversal of the set of all essential separators: the set of all internal nodes. Minimal essential separators and transversal for tree networks are illustrated in Figure 15.

Core-periphery networks

Nodes are divided in two disjoint sets: the core and the periphery. Each node of the periphery is connected to exactly one node of the core, while the nodes of the core are




FIGURE 16. Separators and transversals in core-periphery networks ($n = 12$).

connected with periphery nodes and the core constitutes a clique. Essential separators are subsets of the core. There is a unique transversal: the set of all core nodes. Minimal essential separators and transversals for core-periphery networks are illustrated in Figure 16.

APPENDIX D: ORDER OF MOVES AND NATURE OF CONFLICT

This section explores the role of sequential choice and perfect defense.

Simultaneous moves

Consider a variant of the model studied in the paper where the players make their choice simultaneously. In this case the set of strategies of the defender remains unchanged. A pure strategy of the adversary is now a set of nodes, $X \subseteq N$, chosen to attack. It is important to note that the timing of moves does not affect Lemma 1, which remains unchanged. Suppose that the cost of attack is high. Any strategy, X , in the support of the equilibrium strategy of the adversary must be an individually rational essential separator, i.e., $X \in \mathcal{E}(g, c_A)$. Similarly, any strategy, Δ , in the support of the equilibrium strategy of the defender must be a minimum transversal of the set of essential separators it blocks, $\mathcal{D}(\Delta, \mathcal{E}(g, c_A))$, in $\mathcal{E}(g, c_A)$.

The second observation is that, depending on the network, the players may use pure or mixed strategies in equilibrium. This is a departure from our existing results, where equilibrium always exists in pure strategies. But note that the use of mixed strategies is sensitive to the network. In particular, if the network is such that one unit of defense is sufficient to block all the individually rational essential separators of the adversary, then in equilibrium both players use pure strategies and equilibrium outcomes are the same as in the sequential model studied in the paper. When $\tau(\mathcal{E}(g, c_A)) > 1$, the defender may choose to block more individually rational essential separators by mixing across several transversals.

The model of conflict

We have assumed perfect defense. A more natural way to proceed would be to suppose that the number of resources assigned by each player to a node determines the probability of winning/losing the node. Following Tullock (1980), suppose that the probability of successfully attacking the node is given by a *contest success function* (CSF)

$$\pi(a, d) = \begin{cases} 0 & \text{if } a = 0 \\ \frac{d^\gamma}{a^\gamma + d^\gamma} & \text{otherwise,} \end{cases}$$

where $\gamma \in \mathbb{R}_+$, and a and d are resources assigned by the adversary and defender, respectively. The probability of successfully defending the node is $\pi(d, a) = 1 - \pi(a, d)$.¹¹

A strategy of the defender is a vector $\mathbf{d} \in \mathbb{N}^N$ such that d_i is the number of defense resources assigned to node i . A strategy of the adversary is a function $X : \mathbb{N}^N \rightarrow \mathbb{N}^N$ such that, given vector of defense allocation \mathbf{d} , it maps to a vector of attack allocation $\mathbf{a} = X(\mathbf{d})$ such that a_i is the number of attack resources assigned to node i . We will call the set of nodes that receive a positive number of defense resources the *defended nodes* and the set of nodes that receive a positive number of attack resources the *attacked nodes*. Given defense and attack allocations, (\mathbf{d}, \mathbf{a}) , the probability that set $M \subseteq N$ of nodes is won by the adversary and removed from g is

$$w(M|\mathbf{a}, \mathbf{d}) = \prod_{j \in M} \pi(a_j, d_j).$$

The expected payoffs to the defender and the adversary from defense and attack allocations, (\mathbf{a}, \mathbf{d}) , are

$$\begin{aligned} \Pi^A(\mathbf{a}, \mathbf{d}|g, c_A) &= - \sum_{M \subseteq N} w(M|\mathbf{a}, \mathbf{d})(1 - w(N \setminus M|\mathbf{a}, \mathbf{d}))\Phi(g - M) - c_A \sum_{j \in N} a_j \\ \Pi^D(\mathbf{a}, \mathbf{d}|g, c_D) &= \sum_{M \subseteq N} w(M|\mathbf{a}, \mathbf{d})(1 - w(N \setminus M|\mathbf{a}, \mathbf{d}))\Phi(g - M) - c_D \sum_{j \in N} d_j. \end{aligned}$$

Lemma 1 still obtains. The set of attacked nodes can be decomposed into an essential separator and a reducing attack. In what follows we restrict attention to high costs of attack and we focus on the benchmark model of linear contests: $\gamma = 1$. The main point we wish to make is that with Tullock contests, optimal defense will extend beyond minimal transversals and may cover multiple nodes in the same separator.

Consider an interlinked star with two core nodes: 1, 2, and $n - 2$ periphery nodes ($n \geq 4$). Suppose that the cost of attack is high, $c_A > \Delta f(n - 1)$. The unique essential separator of g is the set of core nodes, $\{1, 2\}$. Let a_1, a_2 be the amount of resources assigned by the adversary to the two core nodes and let d_1, d_2 be the defense resources assigned by the defender to the two core nodes. Expected payoff to the adversary from

¹¹The perfect defense model studied in the paper can be seen as a limiting case of the general contest model: the probability of successful attack is given by $\alpha a^\gamma / (\delta d^\gamma + \alpha a^\gamma)$ with $\alpha = 1$ and $\delta \rightarrow +\infty$.

assignment (a_1, a_2, d_1, d_2) is

$$\begin{aligned} \Pi^A(\mathbf{d}, \mathbf{a}|g, c_A) &= -\pi(a_1, d_2)\pi(a_2, d_2)(n - 2)f(1) \\ &\quad - (\pi(a_1, d_1) + \pi(d_2, a_2) - 2\pi(d_1, a_1)\pi(a_2, d_2))f(n - 1) \\ &\quad - (1 - \pi(a_1, d_1) - \pi(d_2, a_2) + \pi(d_1, a_1)\pi(a_2, d_2))f(n) \\ &\quad - c_A(a_1 + a_2) \\ &= -f(n) + \pi(a_1, d_1)\pi(a_2, d_2)V_1(n) \\ &\quad + (\pi(a_1, d_1) + \pi(a_2, d_2) - \pi(a_1, d_1)\pi(a_2, d_2))V_2(n) \\ &\quad - c_A(a_1 + a_2), \end{aligned}$$

where $V_1(n) = f(n - 1) - (n - 2)f(1)$ and $V_2(n) = f(n) - f(n - 1)$. Notice that $V_2(n)$ is the gain from removing the first node of the core, and $V_1(n)$ is the gain from removing the second node of the core. Since the cost of attack is high, $V_2(n) < c_A$. Hence if $V_1(n) \leq V_2(n)$, then it is not profitable for the adversary to attack, and both players assign no resources to the nodes in equilibrium. Consider now the more interesting case where $V_1(n) > V_2(n)$.

The expected payoff to the defender is

$$\begin{aligned} \Pi^D(\mathbf{d}, \mathbf{a}|g, c_A) &= f(n) - \pi(a_1, d_1)\pi(a_2, d_2)V_1(n) \\ &\quad - (\pi(a_1, d_1) + \pi(a_2, d_2) - \pi(a_1, d_1)\pi(a_2, d_2))V_2(n) \\ &\quad - c_D(d_1 + d_2). \end{aligned}$$

The defender chooses (d_1, d_2) to maximize his expected payoff subject to the constraints that $d_1, d_2 \geq 0$ and that the adversary chooses (a_1, a_2) to maximize his expected payoff subject to $a_1, a_2 \geq 0$.

It is simpler to begin with the case where the defender is given $2d \geq 0$ defense resources and the adversary is given $2a \geq 0$ attack resources. This turns the optimization problem above into a zero-sum bilevel optimization problem, where the defender chooses an allocation of $2d$ to maximize

$$\pi(a_1, d_1)\pi(a_2, d_2)V_1(n) + (\pi(a_1, d_1) + \pi(a_2, d_2) - \pi(a_1, d_1)\pi(a_2, d_2))V_2(n).$$

It is possible to show that the partition (d, d) is a maximizer of both $\pi(a_1, d_1)\pi(a_2, d_2)V_1(n)$ and $(\pi(a_1, d_1) + \pi(a_2, d_2) - \pi(a_1, d_1)\pi(a_2, d_2))V_2(n)$, and hence of the whole expression above. In response the adversary chooses the partition (a, a) . Thus (d, d) and (a, a) are the equilibrium defense and the attack strategies as well.

When both players distribute their resources evenly, the payoff to the adversary is

$$\Pi^A(\mathbf{d}, \mathbf{a}|g, c_A) = -f(n) + \pi(a, d)^2V_1(n) + (2\pi(a, d) - \pi(a, d)^2)V_2(n) - 2c_Aa.$$

If $d \geq V_2(n)/c_A$, it is not profitable for the adversary to attack. Thus with sufficiently low ratio c_D/c_A , the defender distributes his resources evenly and the adversary does not attack. Otherwise both players compete, choosing optimal levels of attack and defense resources and distributing them evenly.

REFERENCES

- Acemoglu, Daron, Azarakhsh Malekian, and Asuman Ozdaglar (2013), “Network security and contagion.” NBER Working Paper 19174. [334, 350]
- Alpcan, Tansu and Tamer Başar (2011), *Network Security: A Decision and Game Theoretic Approach*. Cambridge University Press, Cambridge, England. [334, 335]
- Arquilla, John and David Ronfeldt (2001), *Networks and Netwars: The Future of Terror, Crime, and Militancy*. RAND, Santa Monica, California. [332]
- Aspnes, James, Kevin Chang, and Aleksandr Yampolskiy (2006), “Inoculation strategies for victims of viruses and the sum-of-squares partition problem.” *Journal of Computer and System Sciences*, 72, 1077–1093. [334]
- Baccara, Mariagiovanna and Heski Bar-Isaac (2008), “How to organize crime?” *Review of Economic Studies*, 75, 1039–1067. [334]
- Bala, Venkatesh and Sanjeev Goyal (2000), “A noncooperative model of network formation.” *Econometrica*, 68, 1181–1229. [332, 333]
- Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou (2006), “Who’s who in networks. Wanted: The key player.” *Econometrica*, 74, 1403–1417. [333, 368]
- Bier, Vicki M., Santiago Oliveros, and Larry Samuelson (2007), “Choosing what to protect: Strategic defensive allocation against an unknown attacker.” *Journal of Public Economic Theory*, 9, 563–587. [334]
- Bramoullé, Yann and Rachel Kranton (2007), “Public goods in networks.” *Journal of Economic Theory*, 135, 478–494. [334, 350]
- Cordeiro, Diego, Marcin Dziubiński, and Sanjeev Goyal (2014), “Individual security and network design.” In *Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC '14*, 205–206, ACM, New York. [350]
- Choi, Syngjoo, Andrea Galeotti, and Sanjeev Goyal (forthcoming), “Trading in networks: Theory and experiment.” *Journal of European Economic Association*. [333]
- Clark, Derek J. and Kai A. Konrad (2007), “Asymmetric conflict: Weakest link against best shot.” *Journal of Conflict Resolution*, 51, 457–469. [334]
- Cunningham, William H. (1985), “Optimal attack and reinforcement of a network.” *Journal of the ACM*, 32, 549–561. [334]
- DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel (2003), “Persuasion bias, social influence, and unidimensional opinions.” *Quarterly Journal of Economics*, 118, 909–968. [333]
- Department of Homeland Security (2012), “Office of infrastructure protection strategic plan: 2012–2016.” Technical report. [331]
- Dziubiński, Marcin and Sanjeev Goyal (2013), “Network design and defence.” *Games and Economic Behavior*, 79, 30–43. [334]

Elliott, Matthew and Benjamin Golub (2013), A network approach to public goods. Mimeo, University of Cambridge. [333]

Eun, Ho Oh (2010), *Impact Analysis of Natural Disasters on Critical Infrastructure, Associated Industries, and Communities*. Ph.D. thesis, West Lafayette. [332]

Farrell, Joseph and Garth Saloner (1986), "Installed base and compatibility: Innovation, product preannouncements, and predation." *American Economic Review*, 76, 940–955. [332]

Galeotti, Andrea, Sanjeev Goyal, Matthew O. Jackson, Fernando Vega-Redondo, and Leeat Yariv (2010), "Network games." *Review of Economic Studies*, 77, 218–244. [333]

Golub, Benjamin and Matthew O. Jackson (2010), "Naïve learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*, 2, 112–149. [333]

Goyal, Sanjeev (2007), *Connections: An Introduction to the Economics of Networks*. Princeton University Press, Princeton, New Jersey. [333]

Goyal, Sanjeev and Adrien Vigier (2014), "Attack, defence, and contagion in networks." *Review of Economic Studies*, 81, 1518–1542. [334]

Grötschel, Martin, Clyde L. Monma, and Mehtchild Stoer (1995), "Design of survivable networks." In *Network Models* (Michael O. Ball, Thomas L. Magnanti, Clyde L. Monma, and George Nemhauser, eds.), volume 7 of *Handbooks in Operations Research and Management Science*, 617–672, North-Holland, Amsterdam. [334]

Jackson, Matthew O. (2008), *Social and Economic Networks*. Princeton University Press, Princeton, New Jersey. [333]

Jackson, Matthew O. and Asher Wolinsky (1996), "A strategic model of social and economic networks." *Journal of Economic Theory*, 71, 44–74. [332]

Katz, Michael L. and Carl Shapiro (1985), "Network externalities, competition and compatibility." *American Economic Review*, 75, 424–440. [332]

Kliesen, Kevin L. (1994), "The economics of natural disasters." *The Regional Economist*. [332]

Koutsoupas, Elias and Christos Papadimitriou (1999), "Worst-case equilibria." In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, 404–413, Springer-Verlag, Berlin. [349]

Kovenock, Dan and Brian Roberson (2012), "Conflicts with multiple battlefields." In *Oxford Handbook of the Economics of Peace and Conflict* (Michelle Garfinkel and Stergios Skaperdas, eds.), Oxford University Press, Oxford. [334]

Kunreuther, Howard and Geoffrey Heal (2003), "Interdependent security." *Journal of Risk and Uncertainty*, 26, 231–249. [333]

Luft, Gal (2005), "Pipeline sabotage is terrorist's weapon of choice." *Pipeline and Gas Journal*, 232, 42–44. [332]

Menger, Karl (1927), “Zur allgemeinen kurventheorie.” *Fundamenta Mathematicae*, 10, 96–115. [346]

Pandey, Maneesh (2011), “Political agitations affect railway service.” *India Today*. [332]

Smith, Cole J. (2008), “Preface to special issue on Games, Interdiction, and Human Interaction Problems on Networks.” *Networks*, 52, 109–110. [334]

Tambe, Milind (2012), *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, Cambridge. [335]

Tullock, Gordon (1980), “Efficient rent seeking.” In *Toward a Theory of the Rent Seeking Society* (James M. Buchanan, Robert D. Tollison, and Gordon Tullock, eds.), 97–112, Texas A&M University Press, College Station, Texas. [372]

Vega-Redondo, Fernando (2007), *Complex Social Networks*. Econometric Society Monographs. Cambridge University Press, Cambridge. [333]

Zawadowski, Adam (2013), “Entangled financial systems.” *Review of Financial Studies*, 26, 1291–1323. [350]

Zhu, Shanjiang and David Levinson (2011), “Disruptions in transport networks: A review.” In *Network Reliability in Practice* (David Levinson, M. Henry X. Liu, and Michael Bell, eds.), 5–20, Springer-Verlag, New York. [332]

Co-editor Nicola Persico handled this manuscript.

Manuscript received 29 January, 2015; final version accepted 10 January, 2016; available online 21 January, 2016.