Decentralized Defence of a (Directed) Network Structure

Pelliccia, Marco

Defence and Peace Economics

DOI:
10.1080/10242694.2019.1587679

Published: 17/08/2020

Peer reviewed version

Decentralised Defence of a (Directed) Network Structure

Marco Pelliccia*
Bangor Business School, Bangor University

Abstract

We model the decentralised defence choice of agents connected in a directed graph and exposed to an external threat. The network allows players to receive goods from one or more producers through directed paths. Each agent is endowed with a finite and divisible defence resource that can be allocated to their own security or to that of their peers. The external threat is represented by either a random attack on one of the nodes or by an intelligent attacker who aims to maximise the flow-disruption by seeking to destroy one node. We show that under certain conditions a decentralised defence allocation is efficient when we assume the attacker to be strategic: a centralised allocation of defence resources which minimises the flow-disruption coincides with a decentralised equilibrium allocation. On the other hand, when we assume a random attack, the decentralised allocation is likely to diverge from the central planner's allocation.

Keywords: Networks; Network defence, Security.

JEL classification: C72;

*E-mail: m.pelliccia@bangor.ac.uk. I thank Arupratan Daripa for his insightful comments that were essential to the development of this paper. I am also grateful to the two anonymous referees, Emanuela Sciubba, Nizar Allouch, Arina Nikandrova, Christian Ghiglino, Todd Sandler, Ron Smith, Francesco Cerigioni, and many seminar participants for helpful comments. All mistakes are mine.

1 Introduction

A vast literature has extensively studied the characteristics of games known as *Conflicts on Multiple Battlefields* or *Colonel Blotto games*[1]. In these games, one or more *defendants* defend multiple locations by optimally choosing how to allocate defence resources across themselves, while an intelligent *attacker* aims to conquer as many of them as possible. One of the most important results of these models highlights how a centralised defence allocation is usually more efficient than a decentralised one since it can exploit the negative externalities across multiple locations in order to attract the attacker toward the least valuable ones; individual players fail to internalize the cost of their defence allocation and thus over-invest in defensive measures.

More recently, new contributions have analyzed these games in a network setting (Acemoglu *et al.* (2016), Dziubiński & Goyal (2017), Goyal & Vigier (2014) and Cerdeiro *et al.* (2017)). In these models, the payoff of the players is generally tied to a network structure which connects some of them. This has been motivated by the fact that connections and the architecture of social and economic networks impact the decisions of individuals, firms, and countries in various contexts.[2] For example, an agent may find it beneficial to be part of a large connected component since it may grant him access to a relatively larger amount of goods or to multiple destinations. On the other hand, a terrorist group may aim to disrupt a network infrastructure to damage the welfare of a society which depends on it.

Along the same lines, we propose a model of conflicts where a set of players (defendants) is connected by a directed network structure, and a (unique) attacker aims to maximally disrupt the network by attacking one of its nodes/players. Each defendant benefits from being part of the network as it gives him the possibility to receive goods produced by one or more peers. Each defendant is also endowed with a divisible defence resource which can be transferred to other players. The game is sequential: in the first stage the defendants simultaneously allocate their defence resources, while in the second stage the attacker chooses the node to attack.

We analyze two scenarios. In the first, which we call the *Strategic Scenario* (S1), the attacker is strategic and chooses his attack-strategy in order to maximally disrupt the network given the choices of the defendants. In the second, the *Non-Strategic Scenario* (S0), a node is attacked according to a known probability distribution. By comparing the resulting equilibrium de-

[1] See Kovenock & Roberson (2010), Bier (2006) and Sandler & Enders (2004) for surveys and the works by Bier *et al.* (2007), Lapan & Sandler (1993), Sandler *et al.* (2003), Keohane & Zeckhauser (2003), Kunreuther & Heal (2003), and Heal & Kunreuther (2004).

[2] See Jackson *et al.* (2008).

fence allocations in the two scenarios, we remark that the "strategic element" is an important element to guarantee the efficiency of the decentralized equilibrium.

We first show that when the attacker is strategic, nodes would share defence resources proportionally to their *criticality*[3]. More interestingly, we can show that, under certain conditions, the decentralized defence allocation is *efficient*; it coincides with the defence allocation which minimizes the expected network disruption. On the other hand, when the attack is probabilistic, the players do not necessarily receive defence resources from other peers proportionally to their criticality. In particular, decentralized and centralized defence allocations coincide only under a unique probability distribution over the players. These results complete and, to some extent, challenge the existing literature.

The intuition behind the results is as follows. The directed nature of the network creates a topological ordering along each path connecting a player to a producer. This implies that for each player, his survival and that of any other player in the same path who is crucial to connect him to a producer are equally important. On the other hand, all things being equal, a strategic attacker would prefer to eliminate the most critical nodes. Under certain conditions, this will imply that (i) more critical nodes will receive relatively more defence resources from other peers (Proposition 1), and (ii) the interests of players in the same path will be coordinated, thereby aligning the decentralized allocation to a centralized one (Proposition 2). Loosely speaking, this is possible when the nodes producing goods are connected to any other non-producer node and each player is supported by a number of peers proportional to his criticality. In such a case he would receive enough defence resources so that the attacker finds each node an equally attractive target. When this does occur, this coincides with the allocation criteria of a planner. On the other hand, when the attacker is not strategic but probabilistic, each player would allocate defence to a peer only if the peer is essential to him, and proportionally to his probability of being attacked. In other words, the probability of being attacked, and not the players' criticality, is essentially the unique element which affects a player's defence allocation choice. Conversely, a planner would still take into account the criticality of each node. This is the reason why, under random attack, decentralized and centralized defence allocations coincide only under a unique probability distribution over the nodes (Proposition 4).

Dziubiński & Goyal (2017), Acemoglu *et al.* (2016), Goyal & Vigier (2014), and Cerdeiro *et al.* (2017) are among the closest papers to ours. They study a sequential game in which a designer moves first and chooses a defence allocation, and in a second stage the *adversary*

[3]As we will show in the next sections, a node is more critical if by removing it from the network it has relatively larger impact on the utility of the rest of the nodes.

chooses how to allocate attack-resources across the nodes. In Cerdeiro *et al.* (2017), the authors also discuss how the designer could optimally design the network in order to solve possible inefficiencies arising when security choices are decentralized. In particular, the authors show that decentralised security choices could lead to both over and under-investment in security. In all these works, a strategic attacker targets one node in order to minimise the connectivity of the structure.[4] The main differences with our setting are the following. First, our assumption over the value of the network as perceived by its nodes differs. In our setting, a player profits from being part of a component as far as it allows him to be connected to some producers. In the works previously mentioned, the value of a component is function of its size. A direct consequence of this is that, in our setting, a player might not be affected by the elimination of a node in the same component if this node is not essential to connect the player to a producer. Second, the nature of the attack differs. In Cerdeiro *et al.* (2017), an attack might eliminate a node and propagate to other peers *via* existing links, and in Goyal & Vigier (2014) the attacker can navigate the network by successfully eliminating multiple nodes in multiple rounds. In our setting, there is no contagion and the game terminates after an attack is carried out. This implies again that, as long as the target is not crucial to a player to receive goods, its survival does not impact the player's utility. Although we do not study the optimal network design problem, this difference might also have the following intuitive consequences. In Cerdeiro *et al.* (2017), under strategic attack, in order to incentivise individual nodes to not under-invest in security, a central planner might find it optimal to design a dense network which would make the risk of contagion more likely. In our setting, by making a network more dense, a central planner would only (weakly) reduce the number of likely targets, eventually attracting possible attacks uniquely toward the producer(s). Finally, in the works mentioned, to produce security is costly and the amount of defence produced is a strategic choice. As pointed out by the same authors, this assumption is particularly suitable to describe immunization decision problems.[5] Instead, we focus attention on the reallocation of existing security resources between nodes by allowing them to share defence resources. When studying the problem of decentralized defence of a network structure, two types of inefficiencies might arise. On one hand, individual players might choose to produce inefficient levels of security (over or under-investment in defence). This usually might arise when each player fails to internalize the impact of his choice on the rest of his peers and when se-

[4] Variations of the same problem have been studied by Varian (2004) and Aspnes *et al.* (2006).

[5] There are other notable studies about network flow interdiction problems such as Hong (2011), Wood (1993), Washburn & Wood (1995), Reijnierse *et al.* (1996), Kalai & Zemel (1982), Israeli & Wood (2002). There also exists a vast literature in operations research and computer science about network defence, for instance Alpcan & Başar (2010), Smith (2010), and Zhu & Levinson (2012).

curity decisions are strategic complementary. On the other hand, if allowed to share defence resources with other players, for similar reasons, individual nodes might also inefficiently allocate defence resources. The first type has been the focus of the works mentioned above. Here, we concentrate on the second type of inefficiency, and in order to do so, we separate the individual defence-production problem from the resource-sharing one.

The paper is organised as follows. Section 2 introduces some network notation. Section 3 introduces the model. Section 4 presents the main results. In section 5 we discusses the impact of some network modification on the welfare of the players. Section 6 concludes.

2 Network notation and definitions

A directed network $G(N, L)$ is composed by a set of nodes $N = \{1, ..., n\}$ with $n \geq 2$ and a set of directed links L such that $ij \in L$ means that there exists a directed link from i to j nodes. A *path* between two nodes i and j, P_{ij}, is a sequence of nodes $i_1, i_2, ..., i_k$ such that $i_1 = i$ and $i_k = j$, and $ii_2, i_2i_3, ..., i_{k-1}j \in L$. Two nodes are connected if there exists a path between them. A *cycle* is a path P_{ij} where $i = j$. We define the set of *predecessor nodes*, $B_i \subset N$, as the subset of nodes which can reach i by a path. Similarly the set of *successor nodes* or *follower*, $F_i \subset N$, the subset of nodes which can be reached from i through a path. We say that a node i is a *jq-middleman* node if and only if $i \in P_{jq}$ for any P_{jq}, or there are no paths from j to q which do not involve node i. Thus, we say that i is a middleman if and only if i is *jq-middleman* for at least one ordered pair (j, q) of nodes.[6] We define a node i such that $F_i = \varnothing$, or who does not have any followers, a *sink* node. The *out-degree* of a node i, δ_i^+, is the number of links departing from i, while the *in-degree*, δ_i^-, the number of links received by i. A *star graph* is a graph where a central node is connected to the rest of the players which are uniquely connected to him. A *core-periphery graph* is a graph similar to the star graph where a subset of players composes the core and are connected to the rest of peripheral players. A *directed acyclic graph* (or acyclic digraph) is a directed graph with no cycles. With abuse of notation, we indicate $G - i$ the graph obtained from G by removing the node i and any related link. We finally define by $\mathcal{G}$ the set of directed networks and $\mathcal{G}_n$ the set of directed networks of n nodes.

[6]This definition of middleman may coincide with the widely studied betweenness centrality. However, this is not necessary. The betweenness centrality is measured by considering the shortest paths between two nodes, if more than one, while a node is a *jq-middleman* if any path from j to q passes through him. In other words, a *jq-middleman* would necessarily score a positive betweenness centrality level while a node with positive betweenness centrality score may not be a middleman.

3 Model

There is a set of $(n+1)$ players, $M = N \cup \{A\}$, where N is the set of players which we simply call *defendants*, connected in a directed network $G(N, L)$ with $n \geq 2$, and A is a player which we simply call *attacker*. We call a non-empty subset of nodes $O \subseteq N$ the set of *producers*. Each player $s \in O$ produces a quantity $x_s > 0$ of a good, which can travel through the network *via* the existing directed paths starting from s. Which is, if there exists a path P_{si} in the network G, player i receives the quantity x_s produced by s. We say that two producers, s and q, are *homogeneous* if $x_s = x_q$. Later we will define in details the preferences of each player in N.[7] The next definition will be useful for some of our results.

Definition 1. *A directed network G is connected if and only if each producer $s \in O$ is connected to any $i \notin O$.*

Each node is endowed with a unit of a divisible and transferable resource d which we call *defence resource*. We define $D_i = d_{ii} + \sum_{j \neq i} d_{ji}$, the total defence resources owned by i, where d_{ji} indicates the resource transferred by j to i. We assume that d is non-transferrable to third nodes, which is d_{ji} received by i from j cannot be transferred again to $q \neq j$.

3.1 Conflict

We analyse the two following scenarios:

- **Non Strategic Attack (S0):** One node in N is randomly attacked according to a probability distribution over the nodes set $P(i)$ with $i \in N$.
- **Strategic Attack (S1):** One node in N is attacked by an attacker who aims to maximize the expected disruption of the network.

We specify the technology of conflict. We assume that the attacker A always attacks with a finite constant intensity $\beta > 0$. A node i owning D_i total defence, if attacked, survives with probability $\alpha : \mathbb{R}_+ \to [0, 1)$ which is defined by a classic Tullock contest function[8],

$$\alpha(D_i) = \frac{D_i^\gamma}{D_i^\gamma + \beta^\gamma}$$

[7] To exclude trivial cases, we can assume that any producer has strictly positive out-degree and any non-producer strictly positive in-degree.

[8] See Tullock (2001).

with parameter $\gamma \in (0,1]$ and $\beta > 0$ constant intensity of attack. With probability $1-\alpha(D_i)$, the node is destroyed and thus removed from G. The function $\alpha(D_i)$ naturally captures the ability to resist an attack making it proportional to the relative defence ability of the targeted node i and the one of the attacker. Moreover, the restriction imposed on the parameter γ guarantees strict concavity of $\alpha(D)$ for all $D \geq 0$, or diminishing returns to defence.

Define the *network value function* $v_i : \mathcal{G} \times \mathbb{R}_+ \to \mathbb{R}_+$ as

$$v_i(G) = f\left(\sum_{s \in B_i \cap O} x_s\right) \tag{1}$$

where $B_i \cap O$ is the subset of producers who are also predecessors of i, or the producers who are connected to i. The function $f : \mathbb{R} \to \mathbb{R}$ is differentiable, strictly increasing on $\mathbb{R}_{++}$, and $f(0) = 0$. In words, player i benefits from being part of a component proportionally to the quantity of goods they receive from some producer. If they are not connected to producers, then there is no benefit from being part of the network. Thus, we can naturally compute the *network total value function* $V : \mathbb{R}_+ \to \mathbb{R}_+$ simply as

$$V(G) = \sum_{i \in N} v_i(G)$$

We remark that (1) is a generalization of the network value functions considered in Dziubiński & Goyal (2017), Goyal & Vigier (2014), and Cerdeiro *et al.* (2017). If we assume undirected graphs, or a node can reach any other node of the same a component, and each node is also a producer ($O = N$), then the argument of the function f is essentially a multiple of the component's size. On the other hand, (1) might also describe cases where some path is not available, or where some node may be a simple receiver or an intermediary, i.e. where belonging to a network matters as long as it gives access to specific nodes by a path.[9]

Define a node i's *disruption value* as $\tilde{V}_i = V(G) - V(G-i)$. In words, $\tilde{V}_i$ measures the potential impact of removing i from G on the defendants' valuation of the network. It is always positive, since by removing a player i the total value decreases of at least v_i, and it is maximal when the elimination of a node prevents the rest of the players from receiving any good.

Finally, define m_i as the number of nodes depending on i to be connected to at least one

[9]This may well describe the cases of trade networks, or infrastructure networks for example. Few countries own and export natural resources. The value of belonging to the trade network of a natural resource is linked exclusively to the existence of a trade path from the producer to the final country-consumer.

producer, or $m_i = |M_i|$ where $M_i = \{j \in F_i \ : \ v_j(G) - v_j(G-i) > 0\}$. We can now describe the game in more detail.

3.2 Game setup

We consider a two-stage sequential game. In both S0 and S1, in the first stage the nodes simultaneously choose their defence allocation, while in the second stage one of the nodes is attacked. In S0, the target node is randomly picked among N according to a generic distribution $P(i)$ over N, while in S1, the attacker optimally chooses a probability distribution over N given the choices of the defendants in the first stage.

Each defendant $i \in N$ simultaneously chooses a strategy which is a vector $d_i = (d_{i1}, ..., d_{in})$ with $d_{ij} \geq 0$ and $d_{ii} + \sum_{j \neq i} d_{ij} \leq 1$. Thus the strategy space for each i is $S_i = [0,1]^n$ and $S = S_1 \times ... \times S_n$ the set of strategies. A defendant profile is $S_D = (d_1, d_2, ..., d_n)$.

We focus now on the strategic scenario S1. Given S_D, the attacker chooses an attack profile $S_A(S_D) = (\sigma_1, ..., \sigma_n)$, where σ_i is the probability to attack node i, thus $S_A : S \to \mathcal{L}(N)$ where $\mathcal{L}(N)$ is the set of probability distributions over N. When $\sigma_i = 1$, we refer to a pure strategy. To ease the notation, we simply denote S_A the attack strategy $S_A(S_D)$ when it is clear from the context the relevant defence strategy.

Given the strategy profile (S_D, S_A), the expected payoff of a node i is

$$U_i(G, S_D, S_A) = \sum_{j \in N} \sigma_j [\alpha(D_j) v_i(G) + (1 - \alpha(D_j)) v_i(G - j)]$$

In other words, if the attacker attacks and destroys a node j which is critical to i to receive goods from a producer s, player i gets utility $v_i(G-j) < v_i(G)$. On the other hand, if j is not critical and/or j successfully survives the attack, then payoff of i simply reduces to $v_i(G)$. We assume that if i is attacked and removed from G, then $U_i(G - i, S_D, S_A) = 0$.

The expected payoff of attacker A under (S_D, S_A) is described by $\phi : \mathcal{G} \times S \times \mathcal{L}(N) \to \mathbb{R}_+$ as

$$\phi(G, S_D, S_A) = \sum_{i \in N} \sigma_i (1 - \alpha(D_i)) \tilde{V}_i$$

All things being equal, ϕ is maximal when A attacks and destroys a node i such that $V(G - i) = 0$, and lowest when i survives the attack. The minimal disruption given a successful attack implies $\tilde{V}_i = v_i(G)$. In other words, the attacker's expected payoff increases with the chances of winning the conflict and with the expected disruption caused by the elimination of a target node. The disruption is minimal when it is equal to the network value of the node destroyed while maximal when it is equal to the network total value, or, as result of the successful attack, nobody gets access to any good produced by producers.

A strategy profile (S_D^*, S_A^*) is a sub-game perfect Nash equilibria (SPNE) if and only if for every (S_D, S_A),

- $U_i(G, S_D^*, S_A^*) \geq U_i(G, S_D, S_A^*)$ for all $i \in N$, and
- $\phi(G, S_D^*, S_A^*) \geq \phi(G, S_D^*, S_A)$.

We focus on the SPNE of the game.

4 Results

4.1 Strategic Scenario (S1)

We start by assuming the strategic scenario S1. The first result shows that in any SPNE and network G, we expect an equilibrium defence profile which allocates defence resources to the nodes as proportionally to their disruption values.

Proposition 1. *An equilibrium profile (S_D^*, S_A^*) exists and it is such that for any pair i and j attacked with positive probability,*

$$D_i^* = \left(kD_j^{*\gamma} - \beta^\gamma(1-k)\right)^{\frac{1}{\gamma}}$$

with $k \equiv \tilde{V}_i/\tilde{V}_j$, thus i and j are defended proportionally to their disruption values.

Proof: The existence is guaranteed by the fact that in the second stage, S_A^* is always a best response to S_D^*, and in the first stage, the game played by the defendants has at least one NE since S_i is a compact, convex subset of $[0,1]^n$, and $U_i(\cdot)$ is continuous in $(S_1, ..., S_n)$ and quasiconcave in S_i. To see that $D_i^* \geq D_j^*$ for all $i, j \in N$ such that $\sigma_i = \sigma_j > 0$ and $\tilde{V}_i \geq \tilde{V}_j$, observe that if best response of A is to attack both i and j with positive probability, it must be that $\phi(G, S^*_D, \sigma_i = 1) = \phi(G, S^*_D, \sigma_j = 1)$, or

$$(1 - \alpha(D^*_i))\tilde{V}_i = (1 - \alpha(D^*_j))\tilde{V}_j$$

Thus, $D^*_i \geq D^*_j$, with equality holding only in the case $\tilde{V}_i = \tilde{V}_j$. Rearranging, we obtain the expression for D^*_i as stated in the proposition. □

The intuition is fairly simple. The attacker attacks more than one node with positive probability only if these nodes are perceived as equally attractive. This means that if the attacker targets two nodes of different disruption values, it must be that the node with the higher disruption value is also relatively more defended than the other. Moreover, if multiple nodes are equally crucial to a player for receiving goods from a producer, he would find it optimal to transfer resources to them in a way that would make the attacker indifferent to attacking any particular one of them; if the attacker was instead attacking one of these nodes with probability one and the player could divert some of his defence resources to this target, he would profitably do so.

We are going to check if the decentralized equilibrium defence allocation is *efficient* or if it coincides with the allocation chosen by a planner aiming to minimize the expected network disruption.

Consider the following game played by a central planner (CP) against the attacker A. The CP and A sequentially choose a defence allocation and a probability distribution over the nodes, respectively. Which is, in the first stage CP chooses a vector $D = (D_1, ..., D_n)$ with $D_i \geq 0$ for all $i \in N$ and $\sum_i D_i = n$, and, similarly to the previous setting, in the second stage A chooses a distribution over N given D. The expected payoff of the attacker is not changed. The planner's expected payoff is $\pi(G, D, S_A) = -\phi(G, D, S_A)$. We study the SPNE of the game, (D^e, S^e_A). We call an equilibrium defence allocation D^e an *efficient* defence allocation.

Proposition 2. *For any connected G with homogeneous producers, centralized and decentralized defence equilibrium profiles D^e and S^*_D coincide if $m_i > m_j$ for all $i, j \in N$ such that $\tilde{V}_i/\tilde{V}_j > 1$ and $\sigma^e_i = \sigma^e_j > 0$. The condition is also necessary if $\tilde{V}_i/\tilde{V}_j$ is large enough.*

Proof: We first prove the sufficiency part. Consider a centralized equilibrium profile (D^e, S^e_A) such that $\sigma^e_s = \sigma^e_i > 0$ for at least one pair $i, s \in N$. Assume $m_s > m_i$ and $\tilde{V}_s > \tilde{V}_i$. Without loss of generality, consider the case of one producer $s \in O$ and one non-producer $i \notin O$ such that $\sigma_s^e = \sigma_i^e = 1/2$. The argument for more than one producer and in general more targets would be similar. If $\sigma_s^e = \sigma_i^e$, then $(1-\alpha(D_s^e))\tilde{V}_s = (1-\alpha(D_i^e))\tilde{V}_i$. Suppose $S_D^* \neq D^e$. Start from the case $D_i^e > D_i^*$. In the decentralized setting, this implies that A would optimally attack i with probability $\sigma_i^* = 1$. Best response of m_i players depending on i to receive goods from the producer would be to send resources to i up to m_i, thus if $D_i^e > D_i^*$, it must be that $D_i^e = m_i + 1 + \epsilon > m_i + 1 = D_i^*$, for some $\epsilon > 0$. This means that in the centralized setting it is holding

$$(1-\alpha(n-m_i-1-\epsilon))\tilde{V}_s = (1-\alpha(m_i+1+\epsilon))\tilde{V}_i \tag{2}$$

However, for no $\epsilon \geq 0$ this is satisfied when $m_s > m_i$. In particular, in order to get $\sigma_s^e = \sigma_i^e$, less than m_i resources need to be allocated to i. This is due to the fact that $\tilde{V}_q$ increases linearly with m_q while α increases by less than a unit for each unit defence resource added. Thus, if $D_i^e = m_i + 1 + \epsilon$ with $\epsilon > 0$, then $\sigma_s^e = 1$, a contradiction. Consider now the case $D_i^e < D_i^*$. This implies that $\sigma_s^* = 1$. However, since $m_s = n - 1 > m_i$, followers of s can profitably divert resources to s, until $D_i^* = D_i^e$. Thus S_D^* was not a best response, a contradiction. Therefore, it must be that $S_D^* = D^e$, or centralized and decentralized defence equilibrium profiles must coincide.

We now prove the necessity part. Suppose $S_D^* = D^e$ and for a pair $i, j \in N$ such that $\tilde{V}_i > \tilde{V}_j$, we have $\sigma_i^e = \sigma_j^e = 1/2$. Again, the argument for the case of more than two targets is similar. If $S_D^* = D^e$, then $D_i^e = D_i^* > D_j^* = D_j^e$ and $(1-\alpha(D_j^*))\tilde{V}_j = (1-\alpha(D_i^*))\tilde{V}_i$. We can decompose m_i and m_j followers of the two players as $m_i = m + \bar{m}_j$ and $m_j = m + \bar{m}_i$, where $m \geq 0$ are common followers for whom i and j are both critical and $\bar{m}_j$ (respectively $\bar{m}_i$), the nodes for whom i (respectively j) is uniquely critical. Call $k \equiv \tilde{V}_i/\tilde{V}_j$. Observe that, for $k > 1$ and large enough, there exists a unique $\bar{m}_j^* > 0$ which satisfies

$$\frac{1-\alpha(\bar{m}_i+1)}{1-\alpha(m+\bar{m}_j^*+1)} = k$$

or

$$k\alpha(m+\bar{m}_j^*+1) - \alpha(\bar{m}_i+1) - (k-1) = 0$$

Call the left-hand side $g(k, \bar{m}_j^*)$. It is clear that $\partial g(k, \bar{m}_j^*)/\partial k < 0$ and $\partial g(k, \bar{m}_j)/\partial \bar{m}_j > 0$ for all $\bar{m}_j > 0$. Thus, there exists a $k^* > 1$ such that $\bar{m}_j^* > \bar{m}_i$, or $m_i = m + \bar{m}_j^* > m + \bar{m}_i = m_j$. Therefore, when $k \geq k^*$ for any pair $i, j \in N$ such that $\sigma_i^* = \sigma_j^* > 0$, $S_D^* = D^e$ only if it holds $m_i > m_j$. □

Under certain conditions and when the attacker is strategic, the defendants, by following their individual interests, optimally coordinate their actions and allocate defence resources minimizing the expected network disruption. We explain the intuition by means of a simple example. Consider a network of three nodes connected in a line, where a producer sends a good to the other nodes *via* the second node. It is evident that the producer has the highest disruption value, followed by the middleman and the sink node, respectively. Consider a planner owning three units of defence resources. The planner would allocate the resources such that, if possible, the attacker would find it equally profitable to attack any one of the three nodes; any other allocation would attract the attacker toward one of the nodes with probability one, thus making it profitable for the planner to increase the defence of this node. The planner can achieve this only by allocating resources proportionally to the nodes' disruption values. Consider now the decentralized problem where each node is endowed with a unit defence and assume that the middleman has received more defence than in the planner's allocation. Then, the attacker must find it profitable to attack the producer with probability one, and consequently the other nodes would find it profitable to reallocate some of their defence to the producer. A similar argument holds if the node who initially benefited from more resources was either the producer or the last node of the line. In other words, individual players would redistribute resources in order to maximize their chance of receiving goods from the producer, and this problem, under the conditions stated, coincides with minimizing the expected network disruption, which is the goal of the central planner.

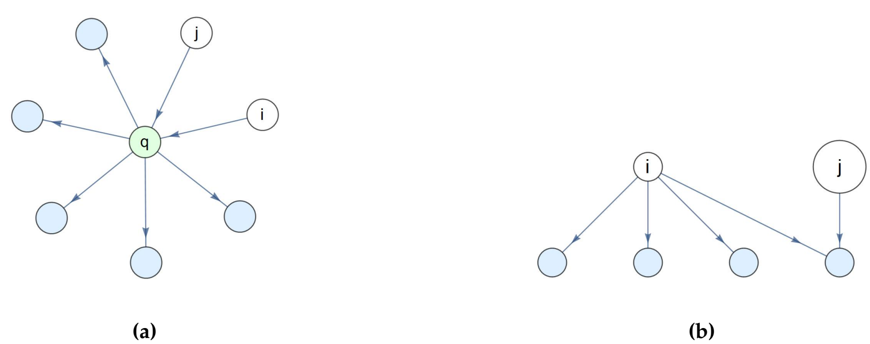

Figure 1: Two examples where conditions for Proposition 2 do not hold and thus decentralized and centralized allocations might differ. In (a), player q is the node with highest disruption value despite $m_q < m_i = m_j$. This might lead to under-protection of q in a decentralized setting since the producers do not strictly benefit from sharing defence with q. In (b), producer j does not reach all the players in N but produce more goods than i. This might lead to under-protection of j in a decentralized setting if $x_j - x_i$ is large enough (in the picture the size of the node indicates the level of production). For instance, a planner might be able to make both the producers equally attractive to the attacker by reallocating optimally n units of defence resources while in a decentralized setting, j would never receive more than 2 units.

More specifically, two conditions guarantee the result. One is related to the producers and one to the followers of targeted nodes. First, if the producers are homogeneous and the network is connected.[10] If this does not hold, we might get an outcome similar to the case of two or more separate components of different size; a planner might still allocate defence such that the attacker is indifferent to attacking two or more players from distinct components while in a decentralized setting this might not happen. Intuitively, players from different components would not strictly benefit from sharing resources between them, thus decentralized and centralized allocation might differ (see Figure 1b).

Second, a node with positive disruption value needs enough followers willing to enhance her defence capability, or the disruption value $\tilde{V}_i$ should be proportional to m_i for each targeted node i. We clarify this point by giving a simple example where this is not the case. Suppose two producers sending goods to a unique middleman q, who in turn is transferring their goods to other $m_q \geq 1$ players (see Figure 1a). In such a case, the disruption value of the middleman, $\tilde{V}_q$, is higher than the disruption values of each individual producer. However, in a decentralized setting, we can only say that q would receive defence with certainty from himself and m_q followers, since the producers do not strictly benefit from sending resources to q. This implies that if in the centralized setting the planner would optimally allocate $D^e_s < 1$ to each producer and $D^e_q > m_q + 1$ to q, there exists at least one equilibrium in the decentralized setting where the defence allocation differs from the planner's allocation. This would not happen if the disruption values of the defendants were proportional to the number of nodes for whom they are critical.

A way to quantify possible inefficiencies arising in the decentralized equilibria is by computing the *price of anarchy* (PoA) (Koutsoupias & Papadimitriou (2009)). This is the ratio between the equilibrium aggregate expected payoff of the defendants in the centralized game and the minimal equilibrium aggregate expected payoff of the defendants in the decentralized setting. Define $W^*(G, S^*_D, S^*_A) = \sum_{i \in N} U_i(G, S^*_D, S^*_A)$ and $W^e(G, D^e, S^e_A) = \sum_{i \in N} U_i(G, D^e, S^e_A)$. For a given network G, the price of anarchy is

$$
\begin{aligned}
PoA(G) &= \frac{W^e(G, D^e, S^e_A)}{\min_{(S^*_D, S^*_A)} W^*(G, S^*_D, S^*_A)} \\
&= \frac{\sum_{i \in N} \sigma^e_i [\alpha(D^e_i)\tilde{V}_i + \sum_{j \in N} v_j(G-i)]}{\min_{(D^*, S^*_A)} \sum_{i \in N} \sigma^*_i [\alpha(D^*_i)\tilde{V}_i + \sum_{j \in N} v_j(G-i)]} \geq 1
\end{aligned}
$$

[10] If the set of producers O is singleton, this condition is trivially satisfied.

which is equal to one whenever the centralized and decentralized allocations coincide. Call $\gamma_i \equiv \alpha(D_i^e)\sum_{j\in N}\tilde{V}_i + \sum_{j\in N} v_j(G-i)$. When $D^e \neq S_D^*$, it must be that there exists a player $i \in N$ such that $D_i^e > D_i^*$ and $\sigma_i^* = 1 > \sigma_i^e$. Hence, the unique decentralized equilibrium profile would imply that m_i followers of i send all their resources to i, or $D_i^* = m_i + 1$, thus the PoA would necessarily be greater than one and equal to

$$PoA(G) = \frac{\sum_{i\in N}\sigma_i^e\gamma_i}{\alpha(m_i+1)\tilde{V}_i + \sum_{j\in N} v_j(G-i)} \tag{3}$$

From (3) we note that, all things being equal, by increasing m_i the inefficiency reduces; if we increase the number of nodes who are willing to defend player i, the defence allocations in the centralized and decentralized setting converge. On the other hand, by increasing the disruption value of node i, holding fixed m_i, the inefficiency increases. For instance, consider again the example in Figure 1a. In a decentralized equilibrium where $D_i^* = D_j^* > 0$ and $\sigma_q^* = 1$, it must be that $D_q^e > D_q^*$. In such a case, if the number of blue nodes were to increase, D_q^e/D_q^* would reduce, thus PoA would decrease. On the contrary, if the number of producers were to increase (white nodes), the disruption value of q would increase while the maximal number of defence resources received by the same would possibly remain unchanged, hence D_q^e/D_q^* would increase and so does the inefficiency as measured by the PoA.

4.2 Non-strategic Scenario (S0)

Consider now S0, or assume that each node $i \in N$ can be attacked according to a probability distribution $P(i)$. As previously, we start by discussing the decentralized setting. Recall M_i, the set of nodes who depend on i to receive goods from some producer.

Proposition 3. *For any G and probability distribution $P(i)$ such that $p_i > 0$ for all $i \in N$, the equilibrium defence profile S_D^* is such that each node i receives resources only from the subset of nodes M_i and proportionally to p_i.*

Proof: Consider the problem faced by a node j connected to a producer s by a unique path P_{sj}. Node j would allocate defence resources to herself and to other nodes in order to maximize the probability of receiving goods from the producer. Since $p_i > 0$ for all $i \in N$, it is never optimal to allocate resources to nodes which are not crucial in order to receive goods from a producer. Since we start by assuming a unique path connecting s to j, any node $q \in P_{sj}$ is essential to j. This means that j would choose an allocation $d_j^* = (d_{jq}^*)$ for all $q \in P_{sj}$ satisfying $\sum_{q \in P_{sj}} d_{jq} = 1$, solving

$$\max_{d_j} P_1 = \sum_{q \in P_{sj}} p_q \alpha(D_q)$$

for a given $P(i)$ and d_i with $i \neq j$. Since P_1 is a sum of strictly concave functions, the optimal allocation d_j^* is unique for each d_i, and such that D_q^* is proportional to p_q. Consider now a node $k \in M_j$. Similarly, player k would choose an allocation d_k^* which, given the usual constraint, maximizes the probability of receiving the good from s, thus solving

$$\begin{aligned} \max_{d_k} P_2 &= p_k \alpha(D_k) + \sum_{q \in P_{sj}} p_q \alpha(D_q) \\ &= P_1 + p_k \alpha(D_k) \end{aligned}$$

Observe that the amount $1 - d_{kk}^*$ optimally allocated by k to the nodes $q \in P_{sj}$ would also maximize P_1 and therefore d_{kq}^* would also be unique given d_i with $i \neq k$, and such that D_q^* is proportional to p_q.

Consider the case of multiple paths connecting producer s to j. Suppose there exists at least another path $P'_{sj} \neq P_{sj}$ connecting s to j. This means that there are at least two nodes $q \in P'_{sj} \cup P_{sj}$ who are not crucial to j to receive the good from s. Without loss of generality, suppose only two alternative paths. Note that the nodes which are essential to j in this case can only be the nodes which are common to both paths, or $z \in P_{sj} \cap P'_{sj}$. Player j's optimal allocation, $d_j^{*'}$, now solves

$$\max_{d_j} P'_1 = \sum_{z \in P_{sj} \cap P'_{sj}} p_z \alpha(D_z)$$

Observe that, since $s \in P_{sj} \cap P'_{sj}$, this set is necessarily non-empty. As previously, this problem admits a solution $d_j^{*'}$ for any d_i since P'_1 is strictly concave, and, given $p_s > 0$, the solution is also unique since $d_{jq}^{*'} = 0$ for all $q \notin P_{sj} \cap P'_{sj}$ and $d_{jz}^{*'}$ is such that D_z^* is proportional to p_z for all $z \in P_{sj} \cap P'_{sj}$.[11] $\square$

[11] It is enough to assume $p_i > 0$ for at least one node i essential to j to be connected to any producer to guarantee $d_{jq}^* = 0$. If there is no such node, sending resources to q would never affect the payoff of j, thus we cannot exclude an equilibrium where $d_{jq}^* > 0$.

Players share defence resources in order to minimize the probability of disruption of crucial paths connecting them to producers. Each node composing a unique path is therefore equally essential to her followers if they are to receive the good. Thus, the only element determining the defence received by one of these nodes when each could be attacked with positive probability is the probability of being attacked, independently of her disruption value. The example in Figure 2 should clarify this point. Player s is the unique producer while player i is a middleman node. Consider beliefs $\{p_s > p_i > 0, p_j = p_q > 0\}$. Since player s is essential to all players and $p_s > p_i$, we expect them to allocate more to s than i ($D_s^* > D_i^*$). Moreover, this is feasible since more players depend on s than i, or $m_s > m_i$, so there will be enough of them willing to satisfy the condition $D_s^* > D_i^*$. Consider instead beliefs $\{p_i > p_s = p_q > p_j > 0\}$. We know that player s and player q would allocate $d_{qq}^* = d_{ss}^* = 1$ to themselves. Player j will then optimally transfer $d_{ji}^* > 0$ to i, and i will allocate $d_{ii}^* = 1$ to himself, so $D_i^* > 1$. Therefore, we obtain $S_D^* = (D_i^* > D_s^* \geq D_q^* > D_j^*)$, which implies individual equilibrium defence capabilities not proportional to the nodes' disruption values.

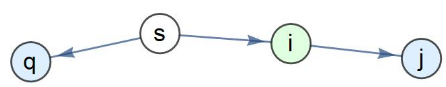

Figure 2: Node s is a producer node while node i is a middleman node.

This example anticipates the following result. Consider again a planner who aims to minimize the expected network disruption. Which is, the planner chooses an allocation $D^e = (D_1^e, ..., D_n^e)$ such that $\sum_{i \in N} D_i = n$ and which minimizes the expected disruption, $\sum_{i \in N} p_i(1 - \alpha(D_i))\tilde{V}_i$ for a given $P(i)$.

Proposition 4. *For any connected G and $P(i)$ such that $p_i > 0$ for all $i \in N$, the centralized and decentralized equilibrium defence allocations coincide if and only if $P(i) = \tilde{P}$, where for all pairs $i, j \in N$, it holds $\tilde{p}_i\tilde{V}_i / \tilde{p}_j\tilde{V}_j = 1$.*

Proof: We prove the proposition for the simple case of $n = 2$ nodes, where one producer s is sending goods to a peer i. The argument for the general case of a network with $n \geq 2$ nodes is the same. Call p_i the probability of i being attacked and $p_s = 1 - p_i$ the probability of s being attacked, and assume both strictly positive. A planner would optimally allocate D_s and

D_i such that the expected disruption $(1-p_i)(1-\alpha(D_s))\tilde{V}_s + p_i(1-\alpha(D_i))\tilde{V}_i$ is minimized given the constraint $D_i + D_s \leq 2$. Clearly the constraint will always bind at the optimum, thus $D_i^e = 2 - D_s^e$ for any D_s^e. Since $\tilde{V}_i = v_i$ and $\tilde{V}_s = v_s + v_i = 2v_i$, the optimal D_s^e would satisfy

$$\frac{\alpha'(D_s^e)}{\alpha'(2-D_s^e)} = \frac{\alpha'(D_s^e)}{\alpha'(D_i^e)} = \left(\frac{p_i}{1-p_i}\right)\frac{1}{2}$$

Consider now the decentralized equilibrium allocation. Given $p_s > 0$, it is easy to see that the producer would always allocate a unit resource to himself, or $d_{ss}^* = 1$. Player i would choose d_{is}^* in order to maximize the probability $(1-p_i)\alpha(1+d_{is}) + p_i\alpha(1-d_{is})$. Therefore, d_{is}^* would satisfy the condition

$$\frac{\alpha'(D_s^*)}{\alpha'(D_i^*)} \leq \left(\frac{p_i}{1-p_i}\right)$$

Note that equality holds if and only if $p_i < 1/2$, while for $p_i \geq 1/2$ we get the corner solution $d_{is}^* = 0$. When $p_i \to 0$, it is trivial to see that $S_D^* = D^e$, or when the producer is attacked with probability close to one, the equilibrium and efficient allocation would coincide to the limit. Observe that by increasing p_i from 0, the difference $D_s^e - D_s^* > 0$ increases, or in the equilibrium allocation s is increasingly under-protected compared to the efficient level. We show that this is true up to $p_i = 1/2$. Call $\tilde{p}$ the probability of A attacking i and such that $D_s^e = 1$. This probability is unique and it is easy to check that $\tilde{p} > 1/2$. We also know that $D_i^* = 1$ for any $p_i \geq 1/2$ since the equilibrium would imply the corner solution $d_{ii}^* = d_{ss}^* = 1$. This means that as far as $p_i < \tilde{p}$, it must be $D_s^e > D_s^*$, while for $p_i \geq \tilde{p}$ it must be that $D_s^e \leq D_s^*$, with equality holding if and only if $p_i = \tilde{p}$. In general, for any G of $n \geq 2$ nodes, there exists a unique distribution $\tilde{P}$ where $\tilde{p}_i > 0$ for all $i \in N$ and such that for all pairs $i, j \in N$,

$$\frac{\tilde{p}_i}{\tilde{p}_j}\frac{\tilde{V}_i}{\tilde{V}_j} = 1 \tag{4}$$

Given $\tilde{P}$, the equilibrium efficient allocation D^e and the decentralized one S_D^* then coincide since $D_i^e = D_i^* = 1$ for all $i \in N$. □

Under S0, when each node can be attacked with some positive probability, decentralized and centralized equilibrium profiles coincide only when the probability distribution over the nodes is $\tilde{P}$ such that $\tilde{p}_i\tilde{V}_i/\tilde{p}_j\tilde{V}_j = 1$ for all pairs $i, j \in N$. This is due to the fact that while the planner would take into account both $P(i)$ and the disruption value of each node when allocating defence resources, individual players would base their allocations purely on $P(i)$.

In particular, we can say that the decentralized equilibrium allocation might be efficient only in two cases. First, when the producers reach all the nodes (network is connected) and the only possible targets are the producers. In such a case, it is intuitive to see that the nodes and the planner all have aligned objectives – to defend the producers. Second, if the attacker attacks each node with positive probability and according to the probability distribution, $\tilde{P}$. In such a case, centralized and decentralized equilibrium defence allocations coincide and imply $D^*_i = D^e_i = 1$ for all $i \in N$. The intuition goes as follows. The planner would allocate resources proportionally to both the nodes' probability of being attacked and their disruption values. When nodes are attacked with probability inversely proportional to their disruption values, individuals' best responses would be to allocate their own resources to themselves. Therefore, decentralized and centralized allocations coincide if and only if the probability distribution over the nodes is such that the planner would optimally allocate equal amount of defence to each node.

Consider for example the network in Figure 2 and assume a random attack such that $p_z = 1/n$ for all $z \in \{s, i, q, j\}$. Equilibrium allocation S^*_D is such that $d^*_{zz} = 1$ for all z, or the nodes do not share defence resources. A central planner would instead allocate resources proportionally to the nodes' disruption value, or $D^e_s > D^e_i > D^e_j = D^e_q$, thus $D^e \neq S^*_D$. On the other hand, suppose $p_s = 1$, or the unique producer is attacked with certainty. Then, $D^*_s = D^e_s = n$, or the decentralized and centralized equilibrium defence allocation coincide. Consider now the probability distribution $\tilde{P} = \{p_s = 0.09, p_i = 0.18, p_j = p_q = 0.36\}$. Then, $S^*_D = D^e$ and such that $D^*_z = 1$ for all $z \in N$. In fact, $\tilde{P}$ is the unique distribution where $p_z > 0$ for all z such that $S^*_D = D^e$.

5 Discussion

5.1 Welfare implications of link-modification

In this section, we discuss how link modifications in G may impact the welfare of the defendants. We assume hereafter the strategic scenario S1.

Given an equilibrium strategy profile (S^*_D, S^*_A), define the set of *potential target* nodes $T \subseteq N$ as $T = \{i \in N : \sigma^*_i > 0\}$. This is the set of nodes who, in equilibrium, are attacked with positive probability by the attacker.

Consider again the welfare function $W(G, S^*_D, S^*_A)$ given the equilibrium profile (S^*_D, S^*_A),

$$W(G, S_D^*, S_A^*) = \sum_{i \in N} \sigma_i [\alpha(D_i^*) V(G) + (1 - \alpha(D_i^*)) V(G - i)]$$

Observe that any network G which in equilibrium maximizes $W(G, S_D^*, S_A^*)$ will also minimizes the attacker's expected payoff $\phi(G, S_D^*, S_A^*)$. This implies that by studying the changes in the attacker's equilibrium expected payoff, we can also infer the relative changes in welfare. Moreover, since the total amount of defence resources in N is finite, by increasing the size of T, we might decrease the total amount of defence owned by a node in T. Therefore, all things being equal, when the number of potential targets increases, the attacker's expected payoff in equilibrium might increase.

To see this more clearly, consider the case of a network with one producer i which is also the unique target, or $T = \{i\}$ (see Figure 3a). In particular, consider the decentralized equilibrium profile where $D_i^* = n$ and $\sigma_i^* = 1$. Let $G' \neq G$ be the network obtained from G by modifying its link structure such that the new set of potential targets, T', is not anymore singleton and one player, node j, becomes a middleman (see Figure 3b). Thus, in equilibrium, $\sigma_i^* = \sigma_j^* > 0$ and $D_i^{*'} < n$ since $D_j^* > 0$. In such a case, the attacker' s expected payoff in G' would necessarily be higher than in G since $\phi(G', S_D^{*'}, S_A^{*'}) = (1 - \alpha(D_s^{*'}))\tilde{V}_s > (1 - \alpha_s(D_s^*))\tilde{V}_s = \phi(G, S_D^*, S_A^*)$. In other words, by increasing the possible targets of A, we also increase the attacker's likelihood to win the conflict.

Definition 2. *A network $G \in \mathcal{G}_n$ is optimal if $W(G, S_D^*, S_A^*) \geq W(G', S_D^*, S_A^*)$ for all $G' \in \mathcal{G}_n$.*

It follows the next result.

Remark 1. *A network G is optimal only if, in equilibrium, is such that $T = O$.*

We provide the intuition here. Consider a network where in equilibrium there exists at least one middleman which is not a producer but which is attacked with positive probability. If there exists a link-modification which can decrease the disruption value of one node in T, it would always weakly improve the expected welfare. Since we can always eliminate middlemen by adding links to the original network, we can always improve the welfare of a network with middlemen. When $T \subseteq O$, there is no link modification which can reduce the size of T, hence an optimal network would necessarily imply an equilibrium where $T \subseteq O$.

(a) G **(b)** G'

Figure 3: In the graph G, the producer i (white node) is also the only potential target node ($T = \{i\}$), thus he will receive defence resources from the rest of the peers ($D_i^* = 5$). In G', the producer i still has maximal disruption value but now $D_i^* = 3.2$ since the middleman j (green) is also critical enough for the rest of the nodes and he will get $D_j^* = 1.8$ ($T' = \{i, j\}$). The expected payoff of the attacker is higher in G' than in G.

We may also ask whether sharing a fixed level of production among multiple nodes is welfare improving. Consider the following three alternative architectures of n nodes and $m = (1, n)$ producers: A core-periphery graph with m core producers, G^{cp} (Figure 4a), a star graph with one central producer, G^s (Figure 4b), and a non-full core-periphery graph, G^{ncp}, where m core producers are connected to each other and to a fraction $(n - m)/m$ of peripheral nodes each (Figure 4c). Assume that the total amount of goods produced by the producers is equal to 1, and when $m > 1$, the total production is equally shared.

Observe that when A attacks with sufficiently high intensity, the unique equilibrium profile in all the three structures is such that A would attack the nodes in the core with positive probability, and the peripheral nodes transfer all their defence resources to these nodes. The welfare levels generated by the three architectures are

$$W(G^s, S_D^*, S_A^*) = \alpha(n) f(1) n$$

$$W(G^{cp}, S_D^*, S_A^*) = \alpha\left(\frac{n}{m}\right) n + \left(1 - \alpha\left(\frac{n}{m}\right)\right) f\left(\frac{m-1}{m}\right)(n-1)$$

$$W(G^{ncp}, S_D^*, S_A^*) = \alpha\left(\frac{n}{m}\right) n + \left(1 - \alpha\left(\frac{n}{m}\right)\right) f\left(\frac{m-1}{m}\right)\left(\frac{m-1}{m}\right) n$$

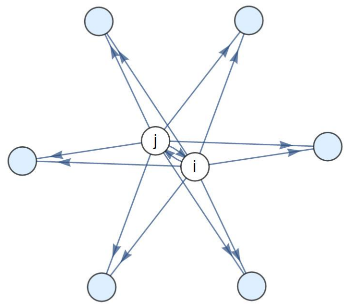

(a) Core-periphery

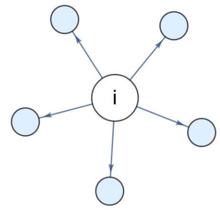

(b) Star

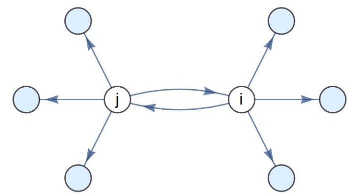

(c) Non-full Core-periphery

Figure 4: The total production is constant and equally shared between producers when more than one. In (a), the removal of one producer would have relative small impact since the rest of $n-1$ nodes could still receive half of the production from the second producer. In (b), the unique producer is maximally defended but his removal gets the highest network disruption. In (c), the removal of one producer gets high disruption although not as high as in (b).

Among the three, the core-periphery graph with m producers is clearly the disruption-minimizing network, but it is also the most "expensive" structure requiring $m(n-1)$ active links. The star graph has the least number of active links, $n-1$, but the disruption in case of failure of the central node is maximal. The non-full core-periphery graph has $(n-2m+m^2)$ active links which makes it less expensive than the core-periphery graph but more expensive than the star graph. In terms of potential disruption, the failure of one of the core nodes in G^{ncp} would prevent the rest of the nodes to get his share of production, $1/m$, and n/m nodes to receive $(m-1)/m$ share produced by the rest of the $m-1$ producers too. This implies that the potential disruption created by a successful attack on one of these nodes in G^{ncp} is between the disruptions expected in G^{cp} and G^s. In terms of resilience, the star graph is the most resilient structure since in equilibrium the unique producer can use all the existing defence resources, while in G^{cp} and G^{ncp}, each producer would own only a fraction n/m of the total defence capability of the population.

It is possible to check that the welfare generated by the core-periphery graph $W(G^{cp}, S_D^*, S_A^*)$ is the highest, while $W(G^{ncp}, S_D^*, S_A^*) \geq W(G^s, S_D^*, S_A^*)$ only for m large enough (see Figure 5a). We provide the intuition. By increasing m, the defence ability of a potential target in G^{cp} decreases as the quantity n/m decreases. However, the disruption value of the potential target also decreases as $n-1$ nodes would possibly lose a smaller share of total production, $1/m$. The latter effect is always stronger than the former when we assume α concave, thus by increasing m, the welfare generated by G^{cp} would increase monotonically. In G^{ncp} each producer is also a middleman for the goods produced by the rest of the producers.

This implies that the disruption value of a potential target would decrease less than in G^{cp} since there would still be a fraction $(n-m)/m$ of peripheral nodes depending on one producer/middleman to receive the whole production from the core. As consequence of this, it is possible that when m is sufficiently small, the impact of a drop in individual's defence ability would be larger than the impact of a reduction in individual production of goods, and consequently, that the welfare generated by G^{ncp} would be smaller than the one expected in G^{s} (see Figure 5a).

We can finally ask how the size of the network, n, might impact the welfare generated by the three architectures considered. In particular, we ask if, for a fixed number of producers m, by increasing n, we expect the differences in welfare to diverge or converge. Let's first consider the graphs G^{cp} and G^{s}. All things being equal, when n increases, a potential target's probability of surviving an attack increases more in G^{cp} than in G^{s} and they both converge to 1. This again is due to the concavity of the α function. Moreover, the disruption created by the failure of a node in the core increases less in G^{cp} than in G^{s}. Hence, all things being equal, by increasing n, the differences in welfare between these two architectures can only increase. Consider now G^{ncp}. When n increases, as before, the probability to survive an attack of a potential target in G^{ncp} and G^{s} converge. On the other hand, the disruption value of a potential target in the two structures also converges since each producer in G^{ncp} becomes a middleman of an increasing number of nodes. Hence, all things being equal, as n increases, we expect the welfare generated by G^{ncp} and G^{s} to converge this time, thus both diverging from the expected welfare in G^{cp} (see Figure 5b).

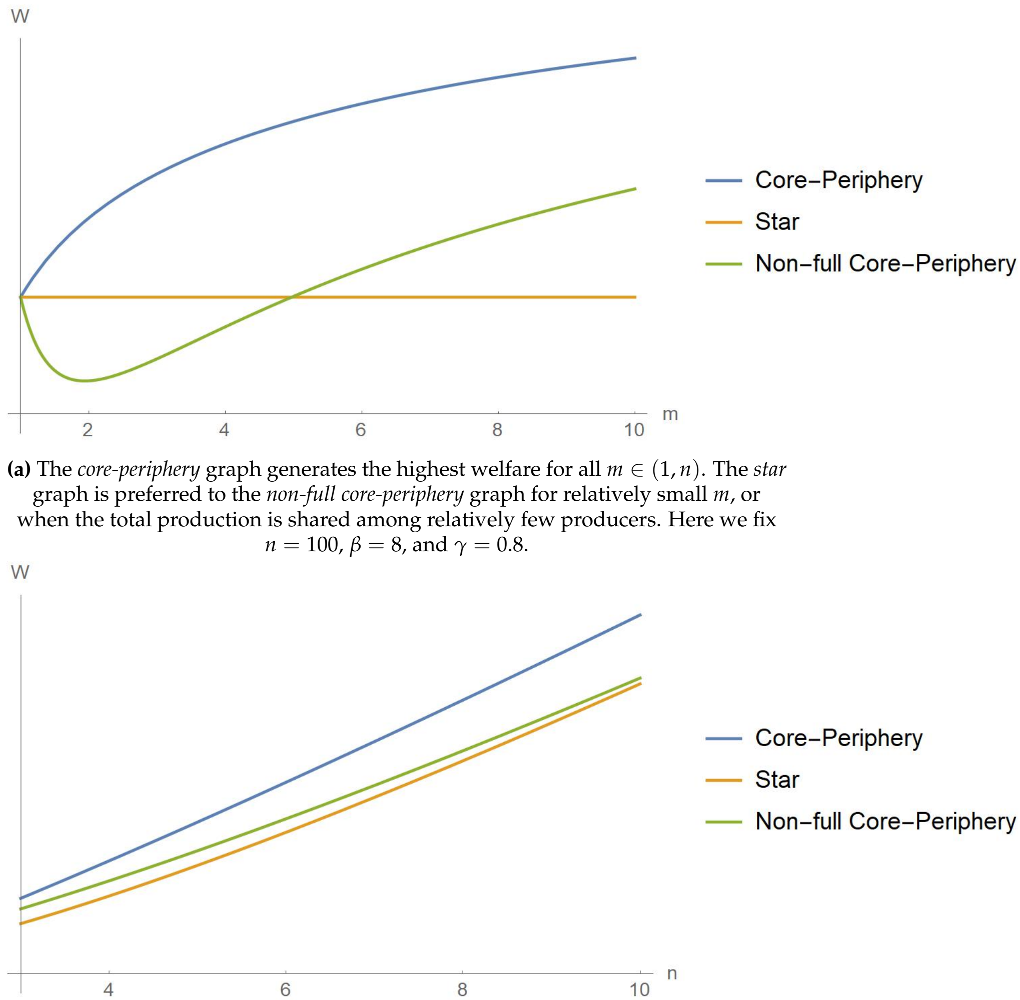

(a) The *core-periphery* graph generates the highest welfare for all $m \in (1, n)$. The *star* graph is preferred to the *non-full core-periphery* graph for relatively small m, or when the total production is shared among relatively few producers. Here we fix $n = 100$, $\beta = 8$, and $\gamma = 0.8$.

(b) The difference between welfare generated by a *core-periphery* structure and the *star* or *non-full core-periphery* architectures increases with the number of peripheral nodes when m is fixed. Here $m = 2$, $\beta = 8$, and $\gamma = 0.8$.

Figure 5

Despite these results being specific to the three structures considered, they can still highlight some important insights. For instance, splitting the production between multiple producers might be welfare improving as long as each individual producer would not become too critical a middleman herself. In other words, when the available defence capability is finite and there are constraints on the number of links we could possibly activate, a planner might face a natural trade-off when designing the optimal network structure: to guarantee high defence ability to each individual target while reducing their individual disruption values.

6 Conclusion

One of the main insights from the literature on games of Conflicts on Multiple Battlefields is that decentralized allocations of defence resources may not be efficient since individual players fail to internalize the negative externalities of their allocation choices and thus over-invest in defensive measures. This has also been confirmed under certain conditions in network settings, or when defendants are connected by a network structure which can be attacked and destroyed by strategic attackers.

We have studied a game from the same family where connected players are endowed with defence units which can be shared between them. We show that if the attacker is strategic (S1), under certain conditions on the network structure, the decentralized allocation of defence resources is efficient, or it coincides with the optimal centralized allocation chosen by a central planner which aims to minimize the expected network disruption. On the other hand, in a non-strategic scenario (S0), we expect an efficient decentralized allocation only under a unique probability distribution over the nodes. This difference is due to the fact that while in S1 players (non-cooperatively) coordinate their actions by taking into account the disruption values of the players in the network, in S0 they do not since the likelihood of an attack on a player is independent of her disruption value.

We also discuss how the network architecture may impact the final welfare of the defendants. Reducing the number of middleman (non producer) players, or players which are crucial to the flow of the goods through the network, is always welfare improving. When the production of goods is shared among multiple producers, core-periphery structures with producers as core players are optimal as their expected disruption is low relative to other architectures. Non-full core-periphery architectures (where each core player is linked to other core players but only to a fraction of peripheral ones) are optimal only when the core is relatively large.

References

Acemoglu, Daron, Malekian, Azarakhsh, & Ozdaglar, Asu. 2016. Network security and contagion. *Journal of Economic Theory*, **166**, 536–585.

Alpcan, Tansu, & Başar, Tamer. 2010. *Network security: A decision and game-theoretic approach*. Cambridge University Press.

Aspnes, James, Chang, Kevin, & Yampolskiy, Aleksandr. 2006. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *Journal of Computer and System Sciences*, **72**(6), 1077–1093.

Bier, Vicki. 2006. Game-theoretic and reliability methods in counterterrorism and security. *Pages 23–40 of: Statistical Methods in Counterterrorism*. Springer.

Bier, Vicki, Oliveros, Santiago, & Samuelson, Larry. 2007. Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*, **9**(4), 563–587.

Cerdeiro, Diego A, Dziubiński, Marcin, & Goyal, Sanjeev. 2017. Individual security, contagion, and network design. *Journal of Economic Theory*, **170**, 182–226.

Dziubiński, Marcin, & Goyal, Sanjeev. 2017. How do you defend a network? *Theoretical Economics*, **12**(1), 331–376.

Goyal, Sanjeev, & Vigier, Adrien. 2014. Attack, Defence, and Contagion in Networks. *The Review of Economic Studies*, **81**(4), 1518–1542.

Heal, Geoffrey, & Kunreuther, Howard. 2004. *Interdependent security: A general model*. Tech. rept. National Bureau of Economic Research.

Hong, Sunghoon. 2011. Strategic network interdiction.

Israeli, Eitan, & Wood, R Kevin. 2002. Shortest-path network interdiction. *Networks*, **40**(2), 97–111.

Jackson, Matthew O, *et al.* 2008. *Social and economic networks*. Vol. 3. Princeton university press Princeton.

Kalai, Ehud, & Zemel, Eitan. 1982. Totally balanced games and games of flow. *Mathematics of Operations Research*, **7**(3), 476–478.

Keohane, Nathaniel O, & Zeckhauser, Richard J. 2003. *The ecology of terror defense*. Springer.

Koutsoupias, Elias, & Papadimitriou, Christos. 2009. Worst-case equilibria. *Computer science review*, **3**(2), 65–69.

Kovenock, Dan, & Roberson, Brian. 2010. Conflicts with multiple battlefields.

Kunreuther, Howard, & Heal, Geoffrey. 2003. Interdependent security. *Journal of risk and uncertainty*, **26**(2-3), 231–249.

Lapan, Harvey E, & Sandler, Todd. 1993. Terrorism and signalling. *European Journal of Political Economy*, **9**(3), 383–397.

Reijnierse, Hans, Maschler, Michael, Potters, Jos, & Tijs, Stef. 1996. Simple flow games. *Games and Economic Behavior*, **16**(2), 238–260.

Sandler, Todd, & Enders, Walter. 2004. An economic perspective on transnational terrorism. *European Journal of Political Economy*, **20**(2), 301–316.

Sandler, Todd, *et al.* 2003. Terrorism & game theory. *Simulation & Gaming*, **34**(3), 319–337.

Smith, J Cole. 2010. Basic interdiction models. *Wiley Encyclopedia of Operations Research and Management Science*.

Tullock, Gordon. 2001. Efficient rent seeking. *Pages 3–16 of: Efficient Rent-Seeking*. Springer.

Varian, Hal. 2004. System reliability and free riding. *Pages 1–15 of: Economics of information security*. Springer.

Washburn, Alan, & Wood, Kevin. 1995. Two-person zero-sum games for network interdiction. *Operations Research*, **43**(2), 243–251.

Wood, R Kevin. 1993. Deterministic network interdiction. *Mathematical and Computer Modelling*, **17**(2), 1–18.

Zhu, Shanjiang, & Levinson, David M. 2012. Disruptions to transportation networks: a review. *Pages 5–20 of: Network reliability in practice*. Springer.