Decentralised Defence of a (Directed) Network Structure

Marco Pelliccia*
Bangor Business School

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 an intelligent attacker who aims to maximise the flow-disruption by seeking to destroy one node. The set of the attacker's potential targets is composed by critical nodes with highest brokerage power and therefore crucial to the system-flow. We show that a decentralised defence allocation is efficient when we assume perfect information: a centralised allocation of defence resources which minimises the flow-disruption coincides with a decentralised allocation. On the other hand, when we assume imperfect information, the decentralised allocation is inefficient.

Keywords: Networks; Network defence, Security.

JEL classification: C72;

*m.pelliccia@bangor.ac.uk

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 multiple *defendants* defend multiple locations by optimally choosing how to allocate defence resources across them, 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.* (2013), Dziubiński & Goyal (2013), Goyal & Vigier (2014) and Cerdeiro *et al.* (2014)). In these models, payoff of the players is generally tied to a network structure which connects part of them. This has been motivated by the fact that connections and the architecture of social and economic networks impact 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.

On 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) strategic 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 of a divisible defence resource which can be transferred to other players. The game is sequential: in the first stage the defendants optimally and 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 target 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 defence

[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).

allocations in the two scenarios, we observe that the "strategic element" may be essential to guarantee the efficiency of the decentralized equilibrium.

We first show that in S1, the equilibrium profile implies that nodes would share and allocate defence resources proportionally to their *criticality*[3]. More interestingly, we can show that the decentralized defence allocation is *efficient*; it coincides with the defence allocation which minimizes the expected network disruption. In S0, this will not necessarily be the case. Nodes which are more critical may end up being less defended than less critical nodes. These results challenge and complete the existing literature. If players benefit from being part of a network structure, then under certain conditions decentralized defence allocations may be efficient.

Dziubiński & Goyal (2013) and Goyal & Vigier (2014) are two of the closest papers to ours. The authors describe a sequential game in which a *designer* moves first and chooses both a network and a defence allocation, and in a second stage the *adversary* chooses how to allocate attack-resources across the nodes. In our setting the defence allocation is decentralised and the attacker can only attack one node with fixed intensity. Furthermore, in both works the network is undirected and its value relies on the size of the connected components. In our model the value of the network as perceived by each node is conditional only on the existence of directed paths from the source(s) to themselves; any modification to the network which does not involve such paths will be payoff irrelevant. In Goyal & Vigier (2014) the authors also study how successful sequential attacks could allow the attacker to "conquer" the entire network. Our game terminates with the attacker's target choice.

Cerdeiro *et al.* (2014) study a similar problem where a designer chooses a network architecture, nodes individually choose their security level, and a strategic attacker targets one node in order to minimise the connectivity of the structure. They show that decentralised security choices could lead to both over and under-investment in security. Our model differs from theirs in two dimensions. Firstly, we assume individual defence resource costs as *sunk*; the quantity of defence resources owned by each node is given and thus not a choice variable. Secondly, each player can optimally reallocate defence resources to themselves or to any other peer in the network.[4]

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

[4]We argue that while in some case it could be reasonable to assume the model setting proposed by the authors (e.g. in the case of a strategic cyber-attack threatening a network of users and the individual choice to invest or not in computer protection), in others it could be more suitable to assume defence investments as a sunk cost which can only be reallocated, *ex-post*, to different locations. This could be the case for example when the resource is a measure of loyalty between agents, or when individual investments in security require time to be fully operative.

Acemoglu *et al.* (2013) analyses the equilibrium and the socially optimal security investments of players connected in generic network structures who are threatened by an external attack. Our setting differs in many aspects. Firstly, our players are able to transfer security resources to other players. Secondly, players' payoffs structure differ substantially.[5]

The paper is organised as follows. Section 2 introduces some network notation. Section 3 introduces the model. Section 4 presents the main results and discusses the impact of some network modification on the welfare of the players. Section 5 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 directed path between them. We define the set of *predecessor nodes*, $S_i \subset N$, as the subset of nodes which can reach i by a directed 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 directed 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 directed 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] 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. With abuse of notation, we indicate $G - i$ the graph obtained from G by removing the node i and any relative link.

[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).

[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 nodes of a directed network $G(N, L)$ with $|N| \geq 2$, and A is a player which we simply call *attacker*. A non-empty subset of nodes $O \subset N$ is composed by *producer* nodes. Each $s \in O$ produces a quantity $x_s > 0$ of a good, which can travel through the network *via* the existing directed links. Later we will define in details the preferences of each player for this good.[7]

Each node is endowed of a unit of a divisible and transferrable 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 A who aims to maximally disrupt the network.

Before characterizing the game, we specify the technology of conflict.

If attacked, a node i survives with probability $\alpha_i(D_i)$ which is a Tullock contest function[8],

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

with parameter $\gamma \in (0, 1)$ and $\beta > 0$ constant intensity of attack. With probability $1 - \alpha_i(D_i)$, the node is destroyed. This simply means that i is removed from G. The function $\alpha_i(D_i)$ naturally captures the ability to resist an attack as proportional to the defence resources owned by the attacked node i. We can now describe the game in more detail.

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

[8] See Tullock (2001).

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** this node is randomly picked among N according to $P(i)$, while in **S1**, the attacker optimally chooses a target node given the choices in the first stage.

Each node $i \in N$ simultaneously chooses a strategy which is vector $d_i = (d_{i1}, ..., d_{in})$ with $d_{ij} > 0$ and $\sum_{j \neq i} d_{ij} = 1$. Thus the strategy space for each i is $S_i = \mathbb{R}^n$ and $S = S_1 \times ... \times S_n$ the set of strategies. A defence profile is $S_D = (d_1, d_2, ..., d_n)$. Given S_D, the attacker chooses one node from N. Which is, a strategy for the attacker A is a function $S_A : \mathbb{R}^n \rightarrow N$.

Define the *network value function* $v_i : G \rightarrow \mathbb{R}_+$ as $v_i(G) = f\left(\sum_{s \in S_i \cap O} x_s\right)$, and the *network total value function* as $V(G) = \sum_{i \in N} f\left(\sum_{s \in S_i \cap O} x_s\right)$. We assume $f'(\cdot) > 0$ and $f(0) = 0$. In other words, each node in G assigns higher value to the network if she can receive more goods by existing directed paths from producer nodes. Each node i has expected utility

$$U_i(G, S_D, S_A) = \alpha_j(D_j)(v_i(G) - v_i(G - j)) + v_i(G - j)$$

with $j \in S_A$. In other words, if the attacker attacks and destroys a node j critical to i to receive x_s from a producer s, i gets utility $U_i(G - j, S_D, S_A) < U_i(G, S_D, S_A)$. If j was not critical to i and/or j resists the attack, then payoff of i stays the same. We assume that if i is removed from G, then $U_i(G - i, S_D, S_A) = 0$.

Define the node i's *total disruption value* $\tilde{V}_i = V(G) - V(G - i)$. Clearly, if $S_i \cap O = \varnothing$ for all i, or there are no nodes which can be reached by a producer node, then $V(G) = 0$. The expected payoff of attacker A given $i \in S_A$ and defence profile S_D is

$$\phi(S_D, S_A, G) = (1 - \alpha_i(D_i))\tilde{V}_i$$

which has highest value, $\phi(\cdot) = (1 - \alpha_i)(D_i)V(G)$, when $V(G - i) = 0$, and lowest, $\phi(\cdot) = 0$, when $V(G - i) = V(G)$. In other words, the attacker's expected payoff increases with the chance of winning the conflict and with the network disruption obtained by removing the targeted node from the network.

A strategy profile (S_D^*, S_A^*) is a pure strategy sequential equilibrium if and only if

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

We study the sub-game perfect Nash equilibria (SPNE) of the game.

4 Results

We start by assuming the strategic scenario S1. The first result shows that in any SPNE and network G, (i) we expect the attacker to attack a node from the producer and middleman sets, and (ii) an equilibrium defence profile which allocates defence resources to the nodes as proportionally to their disruption values.

Proposition 1. *Consider scenario S1. An equilibrium profile (S_D^*, S_A^*) exists and it is such that*

- *A attacks one node from the set of middleman or producer nodes.*
- *For each i, D_i^* is proportional to his disruption value. Which is, for any two nodes i and j such that $\tilde{V}_i > \tilde{V}_j > 0$, it holds*

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

with $k = \tilde{V}_i/\tilde{V}_j$.

We are going to check if the equilibrium defence allocation is *efficient* or if it coincides with the allocation chosen by a central 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 target node respectively. Which is, the 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, similar to the previous setting, the attacker selects a node in N given D. The expected payoff of the attacker is again $\phi(D, S_A, G) = (1 - \alpha_i(D_i))\tilde{V}_i$ with $i \in S_A$ while the CP's expected payoff is simply $\pi(D, i, G) = -\phi(D, i, G)$. We study the sub-game perfect Nash equilibria (D^e, S_A^e). We call an equilibrium defence allocation D^e an *efficient* defence allocation.

Proposition 2. *For any given G, there exists an equilibrium profile (D^e, S_A^*) which coincides with the decentralized equilibrium under strategic scenario S1. Which is, given G and for any equilibrium (D^e, S_A^e), there always exists a correspondent decentralized equilibrium profile (S_D^*, S_A^*) such that $D_i^* = D_i^e$ for all $i \in N$ and $S_A^* = S_A^e$.*

In other words, there exists an equilibrium under scenario S1 which is efficient, or where the nodes optimally coordinate their actions and allocate defence resources such that the expected network disruption is minimized. This is due to two main reasons. First, individual choices of nodes in the same path produce positive externalities on the utility of all the nodes in that path. Second, in absence of any cost to transfer resources between nodes, allocating resources to a non-crucial node does not impact the utility of a sender as far as this would not attract the attacker toward a more crucial-node.

To clarify with a simple example, consider a path which uniquely connects a producer to other $n-1$ nodes. The survival of the producer is crucial to any of these nodes. Similarly, the survival of the second node, direct follower of the producer, is crucial to other $n-2$ follower nodes, and so on. The defendants know that, all things being equal, the attacker value more nodes with higher disruption values. This implies that allocation of resources of more and more distant nodes from the producer will take into account the disruption value of all previous nodes, closer to the producer along the path. Hence, they will adjust accordingly their transfers. For instance, the last node of the path, a sink node, will allocate resources trying to make the attacker indifferent over all the $n-1$ possible targets along the path (the sink node will never be targeted). The previous node, which is one step closer to the producer from the sink node, will allocate resources in the same way. The node two-steps closer to the producer will try to make indifferent the attacker over $n-2$ possible targets, and so on. The resulting allocation will then be inevitably proportional to the disruption values of each node in the path. Moreover, it will be efficient; Any alternative (non-efficient) allocation would make more profitable to the attacker to attack one node in the path over the others and thus it would be optimal to the nodes in the path who depend on this to increase the transfer of defence resources to it.

Suppose now scenario S0, and in particular assume that each node i can be attacked with independent probability $p_i \in [0,1]$ such that $\sum_i p_i = 1$.

Proposition 3. *Consider scenario S0. For any given G and probability distribution $P(i)$, the equilibrium profile S_D^* is such that each node i sends resources to nodes $j \in P_{si}$ with $s \in O$ as proportionally to their probability of being attacked p_j.*

In other words, players share defence resources with other nodes in order to minimize the probability of disruption of paths connecting them to producer nodes. Each node composing a unique path is equally essential to receive the good. However, the final total resources owned by a node may not be proportional to their disruption values. The example in Figure 1 clarifies this point. Node 1 is the unique producer while node 2 is a middleman node. Suppose beliefs $\{p_1 > p_2 > 0, p_3 = p_4 = 0\}$. Since node 1 is essential to all nodes and $p_1 > p_2$, we expect them to allocate more to 1 than 2 ($D_1 > D_2$). Moreover, since $\tilde{V}_1 > \tilde{V}_2$, this is feasible since more nodes depend on 1 than 2, so there will be enough nodes willing to satisfy the condition $D_1 > D_2$. Suppose instead beliefs $\{p_2 > p_1 > 0, p_3 = p_4 = 0\}$. We know that node 1 would always receive at least $d^*_{41} = 1$ and $d_{11} = 1$, so $D_1 \geq 2$. Node 3 will then optimally transfer $d^*_{32} = 1$ to 2, and 2 will allocate $d^*_{22} = 1$ to himself, so $D^*_2 = 2$. Therefore, we obtain $S^*_D = (D^*_1 = D^*_2 = 2, D^*_3 = D^*_4 = 0)$, which implies a level D^*_i not proportional to i's disruption value.

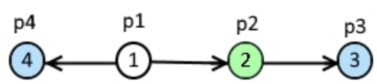

Figure 1: Node 1 is a producer node while node 2 is a middleman node.

This example anticipates the following result. Suppose again a central planner who aims to minimize the expected network disruption. Which is, CP chooses the allocation $D^e = D^e_1, ..., D^e_n)$ which solves

$$\min_D \sum_{i \in N} p_i(1 - \alpha_i(D_i))\tilde{V}_i$$
$$\text{s.t.} \quad \sum_{i \in N} D_i = n$$

Proposition 4. *The CP's optimal allocation differs from the one obtained in S^*_D for any probability distribution $\hat{P}(i)$ such that $p_i > 0$ for all $i \in N$ and acyclic digraphs.*

The proof is trivial so we provide here a simple argument to show how this is always the case. Suppose $p_i > 0$ for all $i \in N$ and a decentralized defence. In any acyclic digraph there always be at least one *sink* node. A sink node, say i, would never receive resources from other nodes, thus $d_{ii} > 0$ for any probability distribution $\hat{P}(i)$ such that $p_i > 0$ for all $i \in N$. However, sink nodes are also nodes such that $\tilde{V}_i = 0$, thus a centralized allocation would imply $D_i^e = 0$, thus $D_i^* \neq D_i^e$. This concludes the argument.

Differently from the result in the strategic scenario S1, under S0 and generic probability distribution $P(i)$, decentralized defence may result in a sub-optimal allocation given our definition of efficient allocation. In this case, nodes in general do not optimally coordinate their action since they do not internalize the negative externality on the rest of their peers. For example, consider again the network in Figure (1) and suppose $p_i = 1/n$ for all i. Equilibrium allocation will imply $D_i^* = 1$ for all i, with $d_{ii}^* = 1$, or each node allocates full resource to himself. A central planner would instead allocate resources proportionally to the nodes' disruption value, i.e. D^e such that $D_1^e > D_2^e$ and $D_3^e = D_4^e = 0$. On the other hand, suppose $p_1 = 1$. Then, $D_1^* = D_1^e = n$, or the decentralized and centralized equilibrium defence allocation coincide.

5 Discussion: Welfare implications of link-modification

We ask how a link modification in G may impact a general utilitarian measure of welfare. In other words, which network architecture does maximize the welfare of the players in N? We assume hereafter strategic scenario S1.

Given the equilibrium strategy profile (S_D^*, S_A^*), define the set of *potential target nodes* $T \subset N$ as $T = \{i \in N : D_i^* > 0\}$. This is the set of nodes who, in equilibrium, own positive defence resources and therefore must be potential targets of the attacker, i.e. A is indifferent over attacking any of them since $\phi(S_D^*, i, G) = \phi(S_D^*, j, G)$ for all $i, j \in T$.

Consider the following Utilitarian Welfare function given target $j \in S_A$,

$$W(D_j, G) = \alpha_j(D_j)\tilde{V}_j + V(G - j)$$

with $j \in S_A$. In other words, the welfare of players in N coincides with the expected total network value given target node j and defence allocation D_j. Note that any network G which in equilibrium maximizes $W(D_j^*, G)$, will also minimizes $\phi(S_D^*, S_A^*, G)$. This implies that by studying the changes in the attacker's equilibrium expected payoff, we can also infer the relative changes in welfare.

We observe that by simply increasing the size of T, more nodes will own positive defence resources. This means that, since the total amount of defence resources in N is finite (equal to n since each node is initially endowed of one unit), sharing it among relatively more nodes would decrease the defence ability of each individual node. Therefore, any increase of size of the potential target-nodes set will increase in equilibrium the attacker's expected payoff.

To see this more clearly, assume G is such that in equilibrium T is a singleton set. From the previous result, this would imply that $i \in T$ gets all resources from the rest of the nodes, thus $D_i^* = n$. Suppose now $G' \neq G$ obtained from G by modifying its link structure and such that T' is not anymore singleton. We can be certain that in the new equilibrium, $D_j^* < n$ for player $j \in T'$ such that $\tilde{V}_j \geq \tilde{V}_z$ for all $z \in T'$. In other words, the node j with highest disruption value in G' will have lower defence capability than the target node i in G. Then, suppose $\tilde{V}_j \geq \tilde{V}_i$. In equilibrium, the expected payoff of A given G will be higher than her expected payoff given G' since $\phi(S_D^{'*}, j, G) = (1 - \alpha_j(D_j^*))\tilde{V}_j > (1 - \alpha_i(D_i^*))\tilde{V}_i = \phi(S_D^*, i, G)$. In general, in G' the attacker's expected payoff in equilibrium can only be higher than under G if and only if $(1 - \alpha_j(D_j^*))\tilde{V}_j \geq (1 - \alpha_i(D_i^*))\tilde{V}_i$, which happens when the disruption value of j node weakly increases and/or the attacker's probability of winning the conflict is not decreased too much. It follows the next result:

Proposition 5. *A network structure maximizes the welfare of the players in* N *only if* $T = O$.

In other words, for a given set of producer nodes O, any architecture maximizing the welfare of players in N will not include middleman nodes. The argument is very intuitive so we will provide here only an informal proof. Any time we increase the number of middlemen, we inevitably increase the disruption value of at least one node (the middleman). This means that the total defence resources would be shared among strictly more nodes. Moreover, the unique nodes who in any structure and equilibrium will receive strictly positive defence resources are the producer nodes. Hence, any time we increase T and make it larger than O, there exists at least one equilibrium where the attacker attacks a producer that now has same disruption value but lower defence resources. Therefore, the attacker will expect higher payoff and thus players will expect lower welfare (see Figure 2).

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

Figure 2: In the graph G, the producer i (white node) is also the only potential target node ($T = \{i\}$). 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 $D_j^* = 1.8$ ($T' = \{i, j\}$). This implies that in G' the expected payoff of the attacker is higher than in G.

We may also ask whether it is welfare improving to share the production among multiple nodes. Consider a star-graph G with one central producer node, say i, and $(n-1)$ peripheral nodes receiving the good from it. The unique equilibrium profile implies that the attacker will attack i and $D_i^* = n$. Consider now the following alternative architecture. Suppose that same quantity produced by the centre of the star-graph is produced by multiple nodes in equal share, say $m \in (1, n)$ producers. Each producer is connected to the rest of $(n-m)$ nodes, thus forming a *core-periphery* structure (see Figure 3a). We may ask under which condition is profitable in terms of welfare to share production among m nodes.

It is easy to see that when each core-node is connected to $n-1$ nodes, we would always increase the expected welfare since $W(D_j, G) < W(D_j, G')$ with $j \in O$, G star-graph and G' the core-periphery graph with $m > 1$ producers in the core. In particular, $W(D_j, G) < W(D_j, G')$ since

$$\alpha_j(n)n < \alpha_j\left(\frac{n}{m}\right)\left(\frac{n-m+1}{m}\right) + \frac{(n-1)(m-1)}{m}$$

for all $m \in (1, n)$.

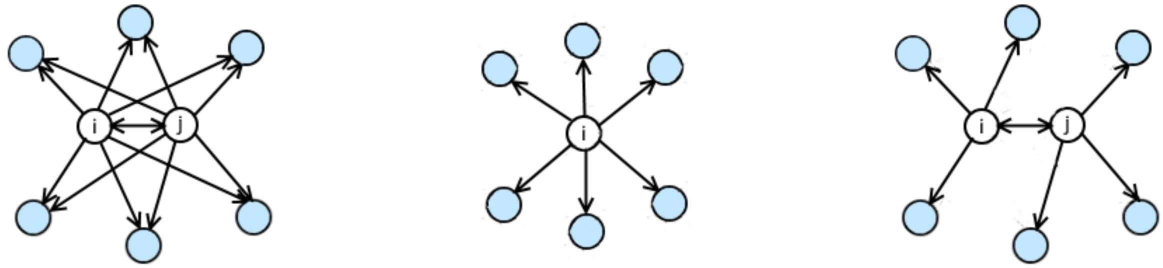

(a) Core-periphery **(b)** Star **(c)** Non-full Core-periphery

Figure 3: 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).

This is fairly intuitive since from G to G', we strictly decrease the disruption value of A's possible targets or G' can only yield higher welfare.

Suppose now that each producer is equally connected only to a fraction of peripheral nodes (see Figure 3c). We call this architecture a *non-full core-periphery* structure. In such case, the conclusion is less clear. In Figure (4) we plot the welfare functions for the star-graph, the core-periphery and the non-full core-periphery with $m > 1$ producers providing the goods for $(n-m)/m$ peripheral nodes each. We can see that there exists a level m^* below which the star-graph with a unique central producer yields higher welfare than the non-full core-periphery graph. In other words, when we share the production among few core nodes and we make each of them middleman for a small group of nodes, they may not receive enough defence resources but they may still have relatively high disruption value.

This point can be particularly relevant if we aim to construct a network architecture which maximizes the welfare of $i \in N$ and we assume a positive marginal cost per-link. We noted that a core-periphery graph with $m > 1$ producers is clearly the most resilient disruption-minimizing network, but it is also the most "expensive" structure requiring $m(n-1)$ active links. This means that, if c is relatively large, also non-full core-periphery architectures may also not necessarily be superior to a star graph ($m = 1$) for mainly two reasons; firstly, the total cost of a non-full core-periphery graph with $m > 1$ producers is $c(n-2m+m^2)$, which increases with m, and is clearly higher than the minimal cost $c(n-1)$ of a star graph. Secondly, for m too small, we have seen that the disruption value of each producer is too high, thus sharing defence resources among m increases the expected payoff of the attacker. Hence, for a given positive cost c, the range of m for which a non-full core-periphery is superior to a star-graph in terms of welfare is even smaller than in absence of cost per-link, i.e. there exists a high level c under which the star-graph yields higher welfare than a non-full core-periphery for all $m > 1$.

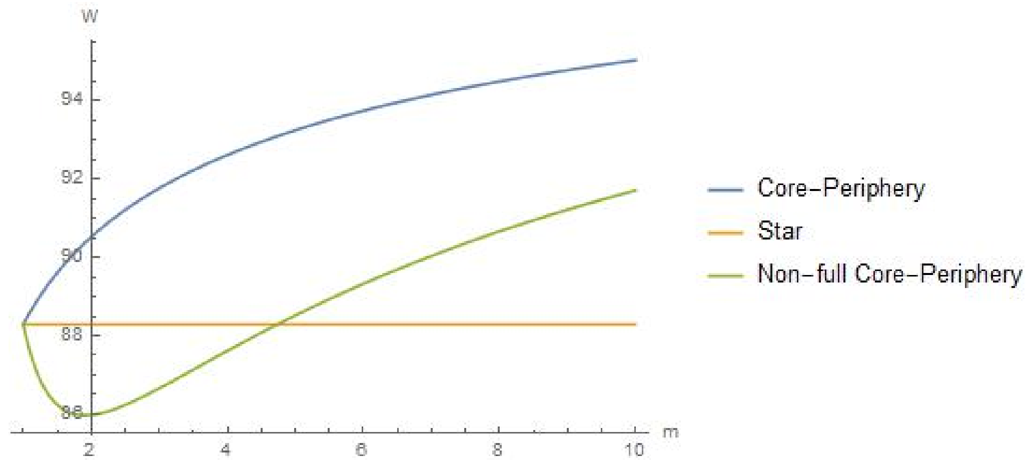

Figure 4: The core-periphery graph gives the highest welfare. The star graph is preferred to the not full core-periphery graph for relatively small m.

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 externality of their allocation 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. Given our setting, we show that in the strategic scenario ($S1$), 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 the non-strategic scenario ($S0$), the decentralized allocation is likely to be not efficient. 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 each node of the network, in $S0$ they do not. In the non-strategic scenario, each node allocates resources in order to minimize the risk of disruption of paths which are crucial to him to receive goods from producer nodes. This implies that nodes fail to take into account the welfare of the rest of their peers in the network. This conclusion may enrich the existing literature by highlighting conditions under which decentralized defence may be efficient.

We finally discuss how the network architecture may impact the final welfare of the defendants. Reducing the number of middleman (non producer) nodes, or nodes which are crucial to the flow of the goods through the network, is always welfare improving. Core-periphery structures with producers as core nodes may be optimal due to their relative low expected disruption but may be expensive to sustain when the core is particularly large and each connection costly. Non-full core-periphery architectures (each core node linked to other core nodes but only to a fraction of peripheral nodes) may be optimal only when core is large enough and cost per-link relatively small.

Proof of the results

Proof of Proposition 1: The existence is guaranteed by the finite and sequential nature of the game and by the perfect information. It is also clear that A will always attack a node from the set of middleman or producer nodes; These are the unique nodes with strictly positive disruption value, thus they are the only nodes which guarantee positive expected payoff if attacked.

We show that D_i^* is proportional to $\tilde{V}_i$. The optimal defence allocation S_D^* will be such that, if possible, A would be indifferent over attacking any of at least two nodes among the middleman and producer ones. This implies that, if A is indifferent over two or more nodes, S_D^* will be such that $\phi(S_D^*, i, G) = \phi(S_D^*, j, G)$ for any i and j among those nodes. Suppose $\tilde{V}_i > \tilde{V}_j$. This implies that $(1 - \alpha_i(D_i^*))\tilde{V}_i = (1 - \alpha_j(D_j^*))\tilde{V}_j$, which, rearranging and replacing $k = \tilde{V}_i/\tilde{V}_j > 1$, gives $D_i^* = \left(kD_j^{*\gamma} - \beta^\gamma(1-k)\right)^{\frac{1}{\gamma}}$. Thus, by increasing $\tilde{V}_i$, all things being equal, D_i^* increases. If there is no S_D which can make A indifferent between two or more nodes, then there exists only one SPNE where A attacks the node with highest disruption value. In such case, it is optimal for the rest of the nodes which depend on it to receive the good to send all their defence resources to the targeted node, i.e. the total defence resources received by the targeted node, say i, will be equal to at least the size of his follower set plus one, $D_i^* \geq |F_i| + 1$. □

Proof of Proposition 2: Suppose $D^e \neq S_D^*$. Then, either $S_A^e \neq S_A^*$, or $S_A^e = S_A^*$. We focus on the second case. If $S_A^e = S_A^*$, then either S_D^* is not a best response, thus S_D^* is not an equilibrium defence profile, or there exists at least one player $i \in N$ such that $v_i(G - j) = v_i(G)$ with $j \in S_A^*$ and $d_{ij}^* < d_{ij}^e$. In the first case, if $j \in S_A^*$ is such that $v_i(G - j) < v_i(G)$ for all $i \in N$, then any S_D^* such that $D_j < D_j^e$ can be improved by i and therefore S_D^* is not an equilibrium defence profile. We can exclude this case. In the second case, observe that for any player i such that $v_i(G - j) = v_i(G)$, sending more resources to i until $D_j^* = D_j^e$ does not affect their expected utility, thus there exists at least one equilibrium where $D_j^* = D_j^e$, or $S_D^* = D^e$, and $S_A^* = S_A^e$. □

Proof of Proposition 3: Suppose a path P_{ij} from a producer i to node j. It is clear that node j will allocate d_{jk}^* to $k \in P_{iz}$ proportionally to p_k in order to maximize the chances to receive the good from i. Call β_1 the j's probability to receive the good from i. This will simply be the product

$$\beta_1 \equiv \prod_q p_q \alpha_q(D_q)$$

with $q \in P_{ij}$. Suppose a node $k \in F_j$, and in particular a direct follower node of j ($jk \in L$). The problem faced by k will be similar to the one of j and in particular k will allocate resources to nodes in the path P_{ik} in order to maximize β_2 computed as

$$\begin{aligned} \beta_2 &\equiv p_k \alpha_k(D_k) \prod_q p_q \alpha_q(D_q) \\ &= p_k \alpha_k(D_k) \beta_1 \end{aligned}$$

Thus, the amount $1 - d^*_{kk}$ will be shared by k across nodes in P_{ij} as proportionally as j has allocated her defence resources to the same. Thus any node belonging to a path will receive resources proportionally to their probability to be attacked. □

References

Acemoglu, Daron, Malekian, Azarakhsh, & Ozdaglar, Asuman. 2013. *Network security and contagion*. Tech. rept. National Bureau of Economic Research.

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

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, Dziubiński, Marcin, & Goyal, Sanjeev. 2014. Individual security and network design. *Pages 205–206 of: Proceedings of the fifteenth ACM conference on Economics and computation*. ACM.

Dziubiński, Marcin, & Goyal, Sanjeev. 2013. Network design and defence. *Games and Economic Behavior*, **79**, 30–43.

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.

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.

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.