Decentralised Defence of a (Directed) Network Structure

Marco Pelliccia*
Department of Economics, Mathematics and Statistics
Birkbeck College, University of London

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 a subset of the set of *middleman nodes* and producers. These are the critical nodes with highest brokerage power in a directed network 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 and under acyclic structures involves no reallocation of defence resources between the nodes. Finally, for a given connected graph, by increasing the link-density we can reduce the set of middleman nodes and thus the number of potential targets. This also decreases the probability of a successful attack.

Keywords: Networks; Network defence, Security.

JEL classification: C69;

*m.pelliccia@mail.bbk.ac.uk

1 Introduction

Institutions, individuals and countries all derive benefit from being part of an interconnected system. A network of pipelines allows the spread of energy resources through different countries; a system of communication channels guarantees a flow of information through members of a community; a supply chain helps to create and sustain competitive advantage in any specific market. In all these cases, as well as in many others, network structures create "value" for the members composing it. A natural question arises: how would they defend their network against a potential external threat?

The types of threats that such networks could face are many and varied. Here, we analyse the case of a strategic *attacker* who targets a single agent of the network. For example, this could be the case of a war-threat against individual countries which are connected by trade-paths.[1] In particular, we model a sequential game where in the first stage each player in the network (simultaneously) decides how to allocate between themselves a defence resource which is infinitely *divisible*, *rival*, and *non-transferable* to third parties. This resource monotonically increases a node's defence ability against an attack. As an example, in the case of a war or terrorist threat, this could either be the share of a country's budget *ex-ante* allocated to domestic and foreign military infrastructures or a measure of *alliance* or *loyalty* between two countries. In the second stage, the attacker chooses her target-node given the defence strategy profile chosen by the players in the first stage. The payoff of the attacker is assumed to increase with the size of the flow-disruption caused by a successful attack.

We study two different scenarios, **S0** and **S1**, which differ in the information available to the *defendants* about the attacker's objectives.[2] Under scenario **S0**, the defendants do not know the attacker's payoff or objectives and therefore, they are assumed to believe that each

[1] Other examples are worth mentioning. As we previously said, a DAG structure could describe a hierarchical structure such as a criminal organization, and the attacker any institution willing to disrupt it. In a supply chain, an attacker could be an incumbent firm strategically targeting a market segment to attack.

[2] Hereafter we define as "defendants" the nodes composing the network to be defended and as "attacker" the external player attacking one of the defendants.

node is attacked with uniform probability. We show that the unique possible equilibrium defence profile implies a inefficient allocation of defence resources irrespective of the network structure considered and, if the network is acyclic, that each node will allocate the entire individual defence resource entirely to herself. In the second setting, **S1**, the defendants have perfect information about the attacker's payoff. In equilibrium, the defendants share resources between themselves taking into account their location in the network structure; the nodes which receive positive defence resources from other peers must belong to the set of the producers and/or *middleman* nodes. The latter are the nodes which have the highest brokerage power among the intermediary nodes in a directed network. Interestingly, when we assume perfect information, the equilibrium profiles turn out to be also *efficient*: a hypothetical central planner aiming to minimize the expected network disruption of a successful attack would allocate the defence resources as in the equilibrium decentralised profile.

One of the main results highlighted by the literature on conflicts pointed out how decentralised defence choices could lead to a *security race* among multiple defenders and consequently to inefficient defence allocations. In particular, this could arise when an increased defence capability of one location may increase the probability that other locations are attacked.[3] We show that in our model this negative externality would not arise. The common value faced by nodes belonging to a connected path from a producer to themselves aligns their incentives towards the resilience of the same path and offsets this type of externality.

In this paper we focus attention on directed networks.[4] This type of structure is observed in systems where one or more producers "send" a good across the network through directed linkages. An example of this could be an arms-trade system; few countries produce arms

[3]See Bier (2006) and Sandler & Enders (2004) for surveys and the works by Lapan & Sandler (1993), Sandler *et al.* (2003), Keohane & Zeckhauser (2003), Kunreuther & Heal (2003), and Heal & Kunreuther (2004) which describe the strategic interactions arising when multiple defenders allocate defence resources across multiple locations.

[4]We also remark the fact that any *undirected* network structure could be reinterpreted in our setting as a directed graph where each connected pair (i, j) is linked by a bilateral edge, or $ij, ji \in L$ with L the set of linkages.

themselves and most therefore rely upon a system of directed trade-paths to receive the product. Supply chains display similar characteristics; (directed) connections between firms from distinct sectors of the economy guarantee the provision of a good to final customers. Finally, any hierarchical structure can be represented by a directed structure: a directed link from i to j can describe a hierarchical order between them which follows the direction of the link.[5]

This paper aims to contribute to the literatures on networks and game theory. To the best of our knowledge, this is the first work studying the strategic decentralized reallocation of defence resources across players connected in a directed network structure. Our model in particular could complement existing ones which study the individual security choices of agents connected in a network structure. Secondly, the directed nature of the graph could effectively describe various networks observed in many social and economic environments where the direction of flows between peers is crucial in defining the value of the same structure to the peers.

The paper is organised as follows. Section 2 introduces the rest of the literature which is relevant to the paper. Section 3 and 4 describes the model setting. Section 5 presents the main results and two examples. Section 6 concludes.

2 Related Literature

The economic literature on conflict is wide. In addition to the works previously mentioned, we should report Major (2002) and in general the class of models called *Colonel Blotto games*. These models describe the zero-sum conflict game between two parties where one party may find it optimal to leave some locations undefended when they are of sufficiently small value and therefore never targeted by the rival. Our model shares this feature, and in particular, due to the fact that our attacker aims to maximise the expected path-disruption, some nodes

[5] More specifically, any hierarchical structure is a directed acyclic graph (DAG) since it allows for a topological ordering over the vertices/nodes and thus implying directed linkages and excluding cycles.

will be naturally "valueless", therefore they will never be attacked and thus defended. Bier *et al.* (2007) describe a game with an attacker and both centralised and decentralised defendants. They show how a centralised defence allocation is 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 defendants fail to internalise the cost of their defence allocation and thus over-invest in defensive measures. In our model, such types of inefficiencies do not arise and a decentralised allocation will be inefficient only under imperfect information. The intuitive reasoning behind this is as follows: Each node has a specific value which is measured in terms of the marginal path-disruption caused by their elimination. Nodes with zero impact will never be targeted by the attacker, thus being undefended would not be enough to attract the attacker to these locations. On the other hand, nodes with positive impacts will be defended proportionally to their value, in order to make the attacker indifferent toward attacking any one of them.

Dziubiński & Goyal (2013), Goyal & Vigier (2014) and Cerdeiro *et al.* (2015) study the impact of the network on the resilience of the structure and on the defence resources allocation. In the first two works, 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 the works the value of the network 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.[6]

Finally, in Cerdeiro *et al.* (2015) the authors study the decentralised defence investment choice of players connected in an undirected network structure. A *designer* chooses a network

[6]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 attack’s target choice.

architecture and a strategic attacker targets one node in order to minimise the connectivity of the structure. In particular, they analyse both the centralised and decentralised choice of security at node-level and 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. 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.

3 Network preliminaries

A *directed graph* $G(N, L)$ is a pair (N, L), where $N = \{1, ..., n\}$ is a finite set of *nodes* and L the finite set of directed edges or links ij connecting a node i to a node j.[7] Being directed, it is not necessarily true that if $ij \in L$ then $ji \in L$.[8] A (directed) *path* in G is a sequence of connected nodes $P_{ij}(G) = \{i_k, i_{k+1}..., i_m\}$ of order $|P_{ij}(G)| \geq 2$, where $i_k = i$, $i_m = j$, $i \neq j$, and $i_k i_{k+1} \in L$ for every $k = 1, ..., m-1$.[9] We say that a node j is *reachable* from a node i if there is a directed path P_{ij}.[10] A directed graph G is *connected* if the underlying undirected graph

[7]With abuse of notation, throughout the paper we will use interchangeably the terms *network* for graph and *player* or *agent* for node.

[8]We assume no *self-loops*, or $ii \notin L$.

[9]To simplify the analysis we consider only connected graphs: there exists a path connecting any pair of nodes. Moreover, in the following analysis we will consider only *simple* paths, or paths where each node appears at most once. To simplify the notation we drop the term "simple".

[10]To ease the notation any elements parametrized to a specific graph G is indicated without specifying G with exception of cases where it is important to distinguish between two or more distinct graphs.

obtained by replacing all directed links with undirected ones is a connected graph, while is *strongly connected* if there exist directed paths P_{ij} and P_{ji} for any pair of nodes (i, j). Define $G_i \subseteq G = (N', L')$ where $N' = \{j \in N : P_{ji} \neq \varnothing\} \cup \{i\}$ and $L' = \{jl \in L : j, l \in P_{ji}\}$. In other words, G_i is a connected component composed by i and any node j from whom i is reachable. The set of all directed paths from a source node s to any other player in N is $\mathbb{P}_s$, while $\mathbb{P}_{ij}$ is the set of all the existing paths from i to j. A *directed acyclic graph* (DAG) is a directed graph with no loops. If $P_{ij} \neq \varnothing$, then we call *descendant* nodes of i any $k \in P_{ij} \setminus \{i\}$ and i is an *ancestor* node of k. A *direct descendant* node of i is a descendant node $j \in N : ij \in L$ and a *direct ancestor* node of i is an ancestor node $j \in N : ji \in L$. The sets of descendant and ancestor nodes of i are $S_i(G) = \{k : P_{ik} \neq \varnothing\}$ and $F_i(G) = \{k : P_{ki} \neq \varnothing\}$ respectively, while the sets of direct descendant nodes and direct ancestor nodes of i are respectively $S_i^d(G) = \{k \in S_i(G) : ik \in L\}$ and $F_i^d(G) = \{k \in F_i(G) : ki \in L\}$. The *out-degree* of i is given by $\eta_i^+ = |S_i^d|$, while the *in-degree* of i by $\eta_i^- = |F_i^d|$. The *degree* of i is simply $\eta_i = |\{S_i^d \cup F_i^d\}|$. We define a node i a *source* if $\eta_i^- = 0$ and $\eta_i^+ \geq 1$, a *leaf* or *sink* node if $\eta_i^+ = 0$ and $\eta_i^- \geq 1$, and *intermediary* node if $\eta_i^+, \eta_i^- \geq 1$. We define the set of producer nodes in G by $O(G) \subseteq N$ and we assume that $O(G) \neq \varnothing$. With abuse of notation, we define $G - i = (N \setminus \{i\}, L - i)$ with $L - i = \{jl \in L : j \neq i, i \neq l\}$ as the graph G (not necessarily connected) after the removal of the node i and any link $ij, ji \in L$. Formally Finally, we indicate with $\mathcal{G}(N)$ the set of all directed graphs over the set of nodes N.

4 Model

Consider a finite set of $(n + 1)$ players $M = N \cup \{A\}$, composed by the set of nodes $N = \{1, ..., n\} : |N| \geq 2$, and the (unique) attacker A. The nodes in N are connected by a directed structure $G(N, L)$ and we assume $O(G) \neq \varnothing$, or there exists at least one node in G which is a *producer*. Moreover, necessary condition for $s \in N$ to be a producer is $\eta_s^+ > 0$ (positive out-degree). To rule out the possibility of loops, or being a producer and a receiver of one's own good, we assume that if there exists a loop W_{ss}, then at least two producers must belong to W_{ss}.[11] Each producer $s \in O(G)$ sends a quantity $x_s > 0$ of a good through the network using the existing directed paths. Each i is endowed with a unit of a divisible defence resource, $d = 1$, which can be allocated to herself and/or to any other player in N. Define with $D_i = d_{ii} + \sum_{j \neq i} d_{ji}$, the total resources owned by node i where d_{ij} indicates the resources allocated by j to i. For each player i, it must hold the following constraint $d \geq d_{ii} + \sum_{j \neq i} d_{ij}$. Any resource received from other peers is assumed to be *non-transferrable* to third nodes, i.e. the quantity $d_{ij} > 0$ allocated by i to j cannot be reallocated by j to a third node $k \neq j$.

If attacked, with probability $(1 - \alpha_i(D_i)) \in (0, 1]$ a node i is "destroyed" along with any directed linkages to or from her. $\alpha_i(D_i) : \mathbb{R}_+ \to [0, 1)$ is a Tullock contest function (Tullock (2001)), continuous and concave with respect to its argument, described by

$$\alpha_i(D_i) = \frac{D_i^\gamma}{\beta^\gamma + D_i^\gamma} \in [0, 1)$$

where $\beta > 0$ is a constant parameter defining the attack resources of A and $\gamma \in [0, 1]$ is a technology parameter.

As previously introduced, we consider two distinct scenarios, **S0** and **S1**, described as follows:

S0: Random attack scenario: Player A attacks one of the nodes with equal probability. Parameters β, γ are common knowledge.

S1: Strategic scenario: Player A attacks strategically. The payoff of A and the parameters β, γ are known by the defendants.

The two distinct cases allow us to check the impact of a strategic environment on the optimal individual choices of the defendants. Below, we describe more in detail the game $\Gamma = \langle M, (\succsim$

[11] If there exists a node $i \in N$ with $\eta_i^+ > 0$ and $\eta_i^- = 0$, a source node, then it must be a producer.

$)_{i \in M}, \delta_{i \in M}, S_A\rangle$. We focus on subgame perfect equilibria (SPNE) of the game Γ.

Timing: Two-stage sequential game. Under **S1**, in the first stage the nodes in N simultaneously choose their defence allocations, while in the second stage, given the nodes' choices, the attacker chooses her unique target. Under **S0**, in the first stage the nodes in N simultaneously choose their defence allocations, while in the second stage there is a random attack by A on one of the nodes in N.

Strategies: Each node $i \in N$ chooses a strategy which is a vector of n coordinates $d_i = (d_{i1}, ..., d_{in})$ where $d_{ij} > 0$ and $\sum_{j=1}^{n} d_{ij} = 1$ which characterizes i's defence allocation among the nodes in $G(N, L)$. All the individual strategies of the nodes determine a defence profile S_D which is a set of vectors composed by all the individual vectors d_i for each i. A strategy for the attacker A is a function $a : \mathbb{R}^n \to \Delta(N)$, which for a given defence profile S_D provides a probability distribution $\Delta(N)$ over the nodes in N. Under **S0** we indicate with $\bar{a}$ the random attack, or the uniform distribution over N. When not specified otherwise, we assume a **S1** scenario. A strategy profile is a pair $(S_D, a(S_D))$. The outcome of a strategy profile $(S_D, a(S_D) = i)$ is a residual network $\tilde{G} \in \{G, G - i\}$. The preferences of the players are determined by the relative utilities from the outcomes of the strategy profiles as described below.

Payoff: A node $i \in N$ consumes the quantity x_s sent by the producer $s \in O(G)$ if and only if is reachable from s. In particular, the payoff of each player $i \in N$ is defined by a function $U_i : G \to \mathbb{R}$ given by $U_i(G) \equiv f(X_i)$, where $f(\cdot)$ is a function increasing in its argument, and $X_i(G) \equiv \sum_{s \in G_i} x_s$. We also assume that $f(0) = 0$ and that the payoff of a player destroyed by A is 0.

Define the following utilitarian welfare function $V : G \to \mathbb{R}$, which we call *network value function*, given by

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

We also define a node i's *disruption value* in a network G as $V(G-i)$. Under strategy $a(S_D)$, attacker A attacks gets payoff which depends of the value of the residual network. In particular her expected payoff is

$$\phi(S_D, a(S_D); G) = -\sum_{i \in N} p_i \left(V(G-i) \left(\frac{\beta^\gamma}{\beta^\gamma + D_i^\gamma} \right) + V(G) \left(\frac{D_i^\gamma}{\beta^\gamma + D_i^\gamma} \right) \right)$$

with $p_i \in [0,1]$ and such that $\sum_{i \in N} p_i = 1$ is the A's probability of attacking node $i \in N$ given strategy $a(S_D)$. With abuse of notation we indicate with $\phi(S_D, i; G)$ as the expected payoff when A attacks player i with probability equal to one.

4.1 Efficient allocation and *Middleman* nodes

In order to compare the equilibrium decentralised defence allocation with a centralised one, we introduce a similar problem faced by a central planner which aims to minimise the impact of a graph-disruption by one attacker. Suppose a central planner endowed of $D = n$ defence resources which chooses a strategy δ^e which is a defence allocation defined again as a vector $D^e = (d_i)$ of n coordinates with positive elements d_i which sum at most n. Similarly to the previous analysis, we assume a two-stage game where the central planner (C) moves first and the attacker A moves sequentially. Again, we analyse two scenarios, **S0** and **S1**. In the strategic scenario, the central planner knows the payoff of A and the parameters β, γ while in **S0** C chooses a defence allocation in the first stage given the following random attack. The expected payoff of C from strategy profile $(\delta^e, a(\delta^e))$ is $-\phi(\delta^e, a(\delta^e); G)$. Under **S1**, due to the perfect information environment and the finite extensive form, we expect the zero-sum game to have an equilibrium in which the absolute payoff from the *minmax* strategy of C coincides with the payoff from the *maxmin* strategy of A.

Definition 1. *Consider **S1**. An allocation D^e is **efficient** if it is a minmax strategy, or*

$$D^e \in \arg\min_{D^e} \max_{i \in N} \phi(\delta^e, a(\delta^e); G) \quad \textit{subject to} \quad \sum_{i \in N} d_i = n$$

*Consider **S0**. An allocation $\bar{D}^e$ is **constrained efficient** if it solves the problem*

$$\bar{D}^e \in \arg\min_{D^e} \sum_{i \in N} p_i \phi(\delta^e, i; G) \quad \textit{subject to} \quad \sum_{i \in N} d_i = n$$

with $p_i = 1/n$ for all $i \in N$.

We can also present the following result.

Lemma 1. *The constrained and efficient allocations coincide.*

Solving the constrained optimisation problem in order to minimise the expected payoff of the attacker, we obtain a minimiser vector of defence allocations which is equal to the one defining the minmax strategy; under random attack scenario, given equal probability for each node of being attacked, the defence allocation is proportional to the value $V(G-i)$ of each $i \in N$. Under strategic scenario, the allocation would be equally proportional to the $V(G-i)$ values and also in order to make the attacker indifferent to attack each one of them.

Since the attacker's objectives concern the graph-disruption due to a successful attack, it is important to highlight which node(s) could be more crucial to the structure and whose destruction would cause the greatest structural damage. Intuitively, each producer is important by construction, since she provides a positive quantity of the good. However, there could exist at least equally crucial intermediary nodes, which are not necessarily producers but that guarantee the flow of the same good to third peers.

Given the directed nature of the linkages, we borrow and modify the definition of *critical node* proposed by Sims & Gilles (2014). They define a node as critical in a given directed network, a *middleman*, if he is essential to connect at least two nodes in a given directed path. This measure is particularly suitable to our setting since it takes into account the directions of the links and the existence of alternative "competing" paths to receive the same flow from two nodes. As extensively discussed by the same authors, these feature are not present in the most standard measure of centrality in the literature such as the *betweenness* centrality.

Definition 2. *Let $G(N, L)$ be a directed and $i, j, h \in N$ be three distinct nodes.*

(a) *The node h is a (i, j)-**middleman** if $P_{ij}(G) \neq \emptyset$ and $h \in \mathbb{P}_{ij}$. We denote by $M_{ij}(G)$ the (i, j)-**middleman set** containing (i, j)-middlemen.*

(b) *The (i, j)-**middleman set** for G is the set of all middlemen in G given by*

$$M(G) = \bigcup_{i,j \in N : i \neq j} M_{ij}(G)$$

*Node h is a **middleman** if $h \in M(G)$ and is a **non-middleman** if $h \notin M(G)$.*

Informally, we define any intermediary node $h \in M_{ij}(G)$ a (i, j)-middleman if he belongs to any directed path from a node i to a node j. Clearly, the set $M_{ij}(G)$ does not necessarily have to be a singleton set.

5 Results

We are going to characterise the equilibrium strategy profiles (subgame perfect Nash equilibria) under the two scenarios considered. We start assuming **S1**, and then we study **S0**. Finally we discuss the equilibrium defence profiles in terms of efficiency. Before presenting the results we formally define the set of *potential targets* $T(S_D^*, a^*; G)$ as the set of nodes which in a given equilibrium profile $(S_D^*, a^*(S_D^*))$ have positive probability to be attacked by A. In particular,

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

- $T(S_D^*, a^*; G) \subseteq M(G) \cup O(G)$
- $d_{ij}^* > 0 \Rightarrow j \in M_{si} \vee j \in O(G)$ *for all* $i \in N$ *and* $s \in O(G)$.

The proposition states that an equilibrium exists and implies that any target chosen by the attacker will belong to the set of middleman nodes and/or producers. Moreover, the unique nodes that are expected to receive defence resources also belong to these sets. Under **S1**, the equilibrium profile implies that the defendants will optimally allocate resources in order to make A indifferent to attacking one of the nodes in $T(S_D^*, a^*; G)$, which is a subset of nodes essential to receive units of good from the producer(s). Since each node could have a distinct disruption value, to make A indifferent, the defence allocation will not necessarily be homogeneous across the nodes but proportional to the nodes' disruption values. In particular, define $k = |\phi((S_D^*, i; G)|$ for all $i \in T(S_D^*, a^*; G)$. It must hold for all i

$$D_i^* = \beta \left(\frac{k + V(G-i)}{k + V(G)} \right)^{\frac{1}{\gamma}}$$

where it is clear that D_i^* proportionally increases with the node i's disruption value. Note that an equilibrium profile will also imply that $\sum_{i \in T} D_i = |\cup S_i(G)|$. In words, the attacker will target a subset of nodes with the highest disruption value. Any node outside the set $T(S_D^*, a^*; G)$ will have zero defence resources and thus, the aggregate amount of defence received by the nodes in $T(S_D^*, a^*; G)$ is equal to the size of the union of their follower sets.

Proposition 2. *Consider **S0** and* $O(G) \subset N$. *Then, any equilibrium configuration implies* S_D^* : $D_i^* = 1$ *for all* $i \in N$ *and* $V(G-i) = V(G-q) > V(G-j)$ *for all* $i, q \in T(S_D^*; \bar{a}, G)$ *and* $j \notin T(S_D^*; \bar{a}, G)$. *Moreover, if G is acyclic,* $S_D^* : d_{ii}^* = 1$ *for all* $i \in N$.

The proposition states that the unique equilibrium profile we expect in **S0** implies that each node will get identical defence resource and if the graph is acyclic, this also implies that each node allocates full resources to himself. This result is mainly driven by the concavity of the contest function and by the ordering of the nodes implied by any directed structure. Note that when choosing the optimal allocation, the nodes take into account only their own probability of being disrupted if attacked by A and not the nodes' general disruption values. Indeed, this will be the main difference with the problem faced by the central planner under **S0**.

Having characterised the equilibrium profiles under the two different scenarios, we can check whether these profiles are efficient. In other words, we can state under which conditions a centralised allocation coincides with a decentralised one.

Proposition 3. *Under **S1**, the equilibrium defence profile S_D^* is efficient while under **S0**, S_D^* is never constrained efficient.*

Under **S1**, the objective of each player is aligned with that of the central planner. In particular, the common goal of each node is to ensure that the path essential for receiving a good from a source is as reliable as possible. Since the objectives of the attacker and the network itself are known by each defendant, they can non-cooperatively coordinate to the efficient allocation which is Pareto optimal. In contrast to one of the common results in the literature, we do not observe a *security race* in this setting. This usually arises when individual security choices create negative externalities on the rest of the peers: an increase in my defence makes other peers of my component relatively more attractive to the attacker. This does not happen in our model since the relatively higher defence capacity of third nodes could increase a node's chance of receiving the good by a path passing through them.

The following remark is worth mentioning. The source nodes do not necessarily receive higher protection than the middleman nodes. For example, imagine two sources connected to a common middleman node which is as sole responsibility for conveying their goods to the rest of the network. In such a case, this middleman node will have the largest disruption value and therefore he is expected to receive relatively higher defence resources than the two sources.

5.1 Example

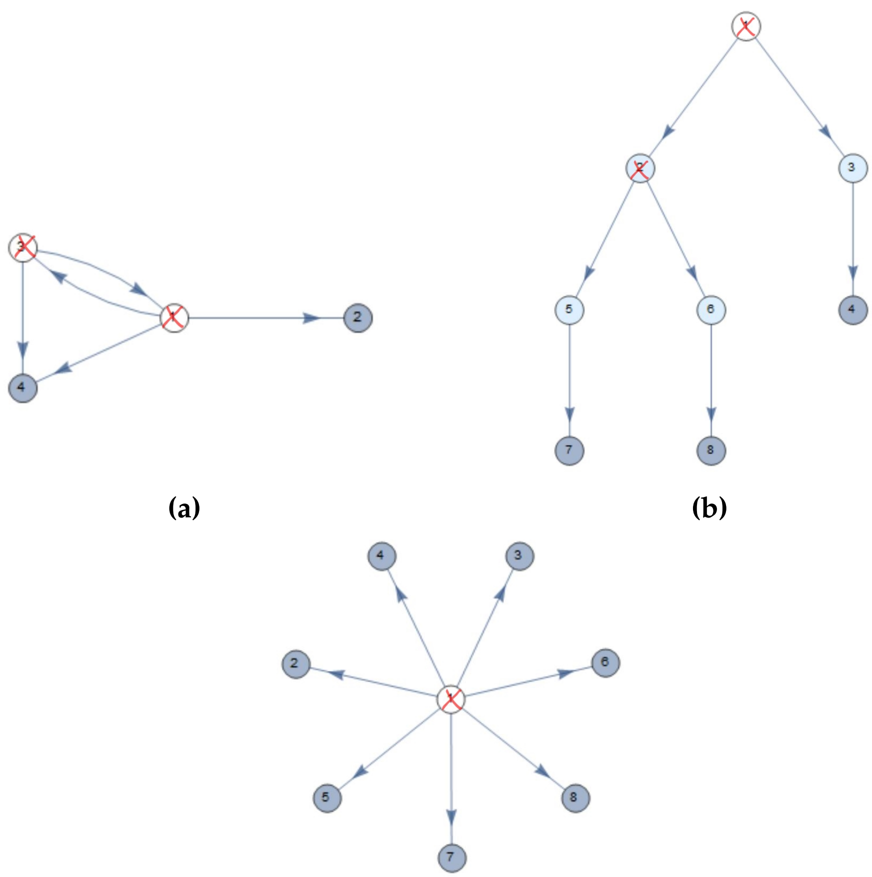

(a) (b)

(c)

Figure 1: Crossed nodes are potential targets (with equal chance of being attacked) under **S1** scenario. The white nodes are the producers, the light blue are intermediary nodes, and the dark blue are sink nodes.

Consider the graph in Figure 1a. The presence of the loop $\{13, 31\}$ *implies that the nodes* $1, 3 \in O(G)$*, i.e. they are both producer and intermediary nodes. The set of middlemen is* $M(G) = \{1, 3\} = O(G)$*. The disruption value of each node is given by* $V(G-1) = 4, V(G-2) = 0, V(G-3) = 3, V(G-4) = 0$*. Assume for simplicity* $\beta = \gamma = 1$*. The equilibrium strategy profile implies*

$D_1^* = 2.3; D_3^* = 1.7; D_2^* = D_4^* = 0$ and $a^* = \{1, 3, p_1 = p_3 = 0.5\}$. Consider the tree graph (DAG) in Figure 1b. The set of producer is singleton, $O(G) = \{1\}$, the set of middleman nodes is $M(G) = \{1, 2, 3, 5, 6\}$, but $T = \{1, 2\}$. In fact, any equilibrium needs to satisfy the following defence and attack profiles: $S_D^* = (D_1^* = 5.4, D_2^* = 2.6, D_i^* = 0 : i = \{3, 4, 5, 6, 7, 8\})$ and $a^* = \{1, 2, p_1 = p_2 = 0.5\}$. Finally, the star graph in Figure 1c implies $S_D^* = (D_1^* = 8, D_i^* = 0 \quad \forall i \neq 1)$ and $a^* = \{1, p_1 = 1\}$.

5.2 Impact of a marginal link

In this section we study the impact of a marginal link in a given connected $G(N, L)$ on the set of potential targets of A. Intuitively, the activation of an extra link in a given connected structure strictly increases the number of existing paths and therefore a node which previously was uncontested may end up to be contested due to the creation of a new path(s): All things being equal, we expect the set $M(G)$ to be weakly smaller for higher network density. Eventually, if each node is directly linked to a source, $M(G) = O(G)$, or the only possible targets of A are the source node(s).

Consider two connected directed graphs $G(N, L)$ and $G'(N, L')$ such that $G' = G + ij$ with $ij \notin L$. Define (S_D^*, a^*) and (S_D', a') the equilibrium profiles in G and G' respectively. Recall that $h \in M(G) \iff h \in \mathbb{P}_{ij}$ for at least one pair $(i, j) \in N$. This directly implies that $M(G') \neq M(G) \iff \exists h \in M(G) : h \notin \mathbb{P}_{ij}(G') \vee \exists h \in M(G') : h \notin \mathbb{P}_{ij}(G)$; In words, the sets of middleman nodes in G and G' differ if and only if there exists at least one node which is a middleman in G but not in G' or vice versa. We can state the first result.

Lemma 2. *Suppose a connected G and $G' = G + ij$ with $ij \notin L$. Then, $|M(G')| \leq |M(G)|$. Moreover, for each $i \in T(S_D^*, a^*; G)$, the probability $\alpha(D_i^*, \beta)$ decreases with $|T(S_D^*, a^*; G)|$.*

In words, for a given connected network structure G, a marginal link ij weakly reduces the size of the set of middleman nodes since a new linkage can only create new paths between a source and other nodes. Moreover, in a connected network an extra link could only either keep the previous middleman nodes uncontested, or make one or more of them contested and therefore not middleman. This happens since the set of all paths existing between two nodes in G is a subset of the set of paths existing between the same nodes in G'. Thus, increasing the density of a given connected structure, we expect the number of middleman nodes, and therefore the set of potential target of A, to weakly decrease. Finally, the attacker could benefit from a greater set of potential targets. This happens since a larger set $T(S_D^*, a^*; G)$ implies a spread of defence resources across more players, and therefore each potential target will be individually weaker in $T(S_D^*, a^*; G)$ than in $T'(S_D^*, a^*; G') \subsetneq T(S_D^*, a^*; G)$.

5.3 Welfare Analysis: impact of higher $|T(S_D^*, a^*; G)|$.

We analyse the impact of increasing the size of the set $T(S_D^*, a^*; G)$ on the sum of the defendants' utilities. Without loss of generality, suppose a graph $G(N, L)$ such that $|N| = n \geq 2$, $O(G) = \{s\}$ and $M(G) = \varnothing$. We know that this setting implies $T = \{s\}$, i.e. the attacker attacks the producer s with probability equal to one. Given the strategy profile (S_D^*, s), the expected network value is

$$W = E[V(\tilde{G})] = \alpha_s(D_s, \beta) n f(x_s) \tag{1}$$

In words, since the attacker will attack the producer s with certainty, the utility of each player in G fully depends on the survival of such node. Consider now an alternative graph $G'(N', L)$ such that again $O(G) = \{s\}$ and $|N'| = n$, but $|M(G')| \geq 1$ and in particular, $|T(S_D^*, a^*; G)| = k > 1$; there exists at least one middleman distinct from the producer such that the potential target set is not singleton. Assume for simplicity that each middleman has the same disruption value. We know that S_D^* will be such that A will be indifferent to attacking any of the nodes in this set. Define the constant probability of being attacked for each player in $T(S_D^*, a^*; G)$ as $p = 1/|T(S_D^*, a^*; G)|$. It is easy to see that the expected welfare now is

$$W' = pnf(x_s)\left[\alpha_s(D_s', \beta) + (k-1)\alpha_j(D_j, \beta)\right] + p(k-1)(1-\alpha_j(D_j,\beta))mf(x_s) \tag{2}$$

where $j \in M(G)$ and $m \equiv n - |S_j|$. Therefore, $W \geq W'$ if and only if

$$\Delta W = n[\alpha_s(D_s, \beta) - p\alpha_s(D_s', \beta) - p(k-1)\alpha_j(D_j, \beta)] - p(k-1)m(1-\alpha_j(D_j,\beta)) \geq 0 \tag{3}$$

Note that

$$\frac{d(\Delta W)}{d\beta} \leq 0 \iff \left(\frac{d\alpha_s(D_s,\beta)}{d\beta} - p\frac{d\alpha_s(D_s',\beta)}{d\beta}\right) \leq \frac{(n-m)}{n}p(k-1)\frac{d\alpha_j(D_j,\beta)}{d\beta}$$

or for m close enough to n, ΔW monotonically decrease with respect to β for some large enough β value. Moreover, for m large enough, if $\beta \to 0$ then $\Delta W > 0$, while if $\beta \to \infty$ then $\Delta W < 0$. Therefore, we can find a value, $\beta^* > 0$, which reduces $\Delta W = 0$. In other words, there exists an attacker's resource level β^* such that for any $\beta < \beta^*$, network G leads to higher welfare than G', while for any $\beta > \beta^*$ the opposite is true.

We remark that this result hold for any graph G with one producer and $k-1$ middleman nodes which are part of the potential targets of A, and which have identical disruption value. In particular, the specific architecture of G does not affect the qualitative result of the analysis.

In the plot of Figure 2, we show as example the ΔW function for G the star graph of $n = 8$ peripheral nodes and one producer in Figure 1c and G' is a tree graph as the one in Figure 1b. In the Appendix section **A2** we propose a similar welfare comparison between a star graph and a very common network structure, the core-periphery graph. There, we aim to study under which conditions it may be beneficial to spread the production evenly among more than one producers. In equilibrium we expect a lower individual defence ability of the potential targets but also a lower impact of a successful attack in terms of network disruption.

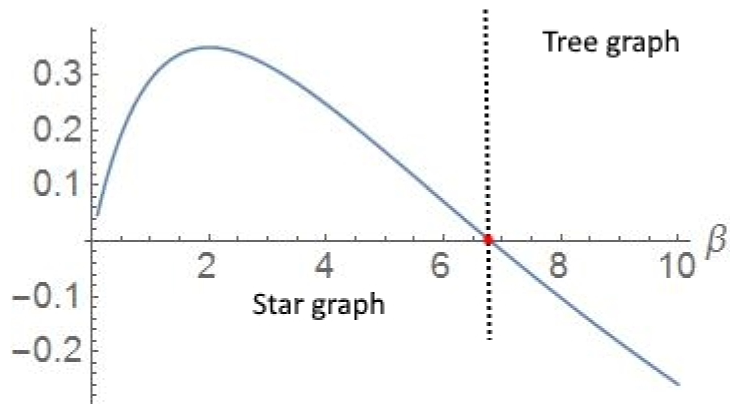

Figure 2: The threshold β level is $\beta^* \approx 6.8$; for $\beta < \beta^*$ the star graph guarantees higher welfare than the tree structure, while the opposite is true for $\beta > \beta^*$.

6 Conclusion

We study the decentralised defence resources allocation by players connected in a directed network structure. The players face the threat of a strategic attacker aiming to maximally disrupt the flow from the producer(s) to the rest of the peers. In particular, we consider two scenarios which differ in the information owned by the defendants. We find that, under perfect information (**S1**), the nodes optimally share resources between them as function of their relative location in the network. In particular, the nodes which receive positive defence resources belong to the subset of middleman nodes and producers. These are the nodes whose location is critical to the flows of goods through the network. On the other hand, under imperfect knowledge (**S0**), players do not necessarily share defence resources and instead if the network structure is acyclic they allocate their entire defence endowment to themselves. Interestingly, under **S1** the decentralised allocation is also efficient: the equilibrium defence profile coincides with that of a central planner aiming to minimise the potential flow disruption.

In the model proposed we consider an attacker which is able to attack only one node and with fixed intensity. Further extensions could relax this assumption allowing the attacker to allocate attack-resources to multiple locations. Moreover, in our setting the network structure is given. Many applications could ask to relax this assumption and analyse the consequences of a strategic rewiring of nodes' linkages. This extension in particular could open interesting new research questions. Intuitively, the nodes could face a trade-off between the marginal benefit from activating an extra link in order to decrease their dependency on middleman nodes and its marginal cost given by a maintenance price per link. Finally, a natural extension of the model could study the problem of a central planner choosing both the optimal defence allocation across the nodes and the architecture of the graph. This could help to measure potential inefficiencies between the decentralised and centralised allocation of defence resources.

A Appendix

A.1 Proofs of the results

Proof of Proposition 1. The existence of an equilibrium profile is guaranteed by the finite and sequential nature of the game and by the perfect information scenario. To see that $d^*_{ij} > 0 \Rightarrow j \in M_{si} \vee j \in O(G)$ with $s \in O(G)$, suppose $d^*_{ij} > 0$ and $j \notin M_{si} \wedge j \notin O(G)$. Then the difference $U_i(G-j) - U_i(G-q) = 0$ for any $D_j \geq 0$ and $q \neq j$, thus d^*_{ij} would be ineffective. We know that for any player $i \in N$ there always exists at least one player q such that $U_i(G-q) - U_i(G-t) > 0$ with $t \neq q$. This means that i could improve her chance of receiving a good by transferring defence resources from j to q, thus $d^*_{ij} = 0$.

We can show that S^*_D is such that $T(S_D, a(S_D); G)$ is as large as possible. Suppose that there exists a defence allocation S_D satisfying the condition stated above and such that $\phi(S_D, i; G) = K > 0$ for all $i \in T$. For simplicity and without loss of generality, suppose $|T| > 1$. It is clear that any deviation from this allocation will imply a singleton set $\tilde{T} \neq T$. Suppose $i \in T$ and a deviation by a node $j : d_{ji} > 0$ which consists in a reduction of defence sent to i, say $d'_{ji} < d_{ji}$, and thus $\phi(\tilde{S}_D, i; G) > K$. Then, the best response of A would be $a^*(\tilde{S}_D) = i$ with probability one, or $\tilde{T} = \{i\}$. Then node j could optimally reallocate resources to i in order to increase i's probability of surviving the attack until we again obtain $\phi(S_D, i; G) = K$, and thus $i \in T$. □

Proof of Proposition 2. We start showing that the profile $S^*_D : d^*_{ii} = 1$ for all $i \in N$, and $a^*(S^*_D)$ such that $V(G-i) = V(G-q) > V(G-j)$ for all $i, q \in T$ and $j \notin T$ is indeed a SPNE. Suppose node i deviates from this profile and thus $d_{ii} < 1$ and $d_{ij} > 0$ for some $j \neq i$. Then, since $d_{jj} = 1$, $D_j > 1 > D_i$. Given the concavity of the Tullock function assumed and the equal chance for each node of being attacked, the deviation is not profitable, thus $d^*_{ii} = 1$.

Suppose G is acyclic. First note that if there exists a source s and a sink node q, then $d_{ss} = d_{qq} = 1$. This is a consequence of the fact that $d^*_{ij} > 0 \Rightarrow j \in F_i(G)$ for all $i \in N,\ j \neq i$.

Therefore, we only need to check if $d_{ii} < 1$ for at least one intermediary node $i \in N$. Suppose an intermediary node i with the sink node $q \in S_i^D$ as direct follower. Then i will also allocate $d_{ii} = 1$ since $d_{qi} = 0$. Moreover, any acyclic structure has a partial order $\geq$ on its nodes, thus for the same argument any intermediary node k, ancestor of i will choose $d_{kk} = 1$. In the case that there are no sink and source nodes, the assumption which excludes the possibility for any producer of receiving his good in a loop allows for a partial order of the nodes, and for the same reason previously stated $d_{ii}^* = 1$.

However, suppose G is cyclic. The equilibrium configuration can only guarantee the more general condition $D_i^* = 1$ for all $i \in N$. To see this, consider a cycle graph of $n \geq 4$ nodes and thus at least two producers. It is easy to see that $d_{ij}^* = 1$ for all i, j such that j is the ancestor node of i is indeed an equilibrium defence profile; any deviation implies that a path from one producer to other nodes is not optimally defended and therefore the deviation cannot be profitable. This concludes the proof. □

Proof of Proposition 3. The case under **S0** is trivial and therefore omitted. Consider the scenario **S1**. Suppose an equilibrium profile (S_D^*, a^*). We know that $|T| \geq 1$, $\sum_{i \in T} D_i = n$, the expected payoff of A is $\phi(S_D^*, a^*(S_D^*); G)$ with $p_i = 1/|T|\ \forall i \in T$, and finally $V(G-i) > V(G-k)$ and $D_i^* > D_k^*\ \forall i \in T,\ k \neq i$. Suppose by contradiction that there exists another equilibrium profile $(\tilde{S}_D, \tilde{a}) \neq (S_D^*, a^*)$, which is efficient. This means that $\phi(\tilde{S}_D, \tilde{a}(\tilde{S}_D); G) < \phi(S_D^*, a^*(S_D^*); G)$. Since in equilibrium it must be that $\phi(S_D^*, i; G) = q\ \forall i \in T$ and $\phi(\tilde{S}_D, j; G) = t\ \forall j \in \tilde{T}$ with q and t positive real numbers, we can rewrite the condition simply as $t < q$. Initially we stated that (S_D^*, a^*) implies $V(G-i) > V(G-k)\ \forall i \in T\ , k \notin T$, thus since $(\tilde{S}_D, \tilde{a})$ is an equilibrium profile, it must be that $\tilde{T} \subseteq T$ or $\tilde{T} \supseteq T$.

We are going to prove that $T = \tilde{T}$, or $\tilde{T} \not\subset T$ and $\tilde{T} \not\supset T$. Suppose $\tilde{T} \subsetneq T$. This means that there exists a player $w \in T$ and $w \notin \tilde{T}$. However, this also implies that $\phi(\tilde{S}_D, q; G) < \phi(\tilde{S}_D, j; G) < q$ where $D_q = 0$ and therefore $q \notin T$, which is a contradiction. Thus, $\tilde{T} \not\subsetneq T$. Suppose $\tilde{T} \supsetneq T$. Then there exists some node who belongs to $\tilde{T}$ but not to T. Define the nodes w such that $w \in \tilde{T} \setminus T$ and nodes k such that $k \in \tilde{T} \cap T$. We know that for any node $j \in \tilde{T}$, it must be that $\phi(\tilde{S}_D, j; G) = t$ and $D_j > 0$. Thus, $\phi(\tilde{S}_D, w; G) > \phi(\tilde{S}_D, j; G)$ where $D_w = 0$ and $D_j > 0$. The initial equilibrium profile implies that any node w is such that $\phi(\tilde{S}_D, w; G) < \phi(\tilde{S}_D, k; G)$ where $D_w = 0$ and $D - k > 0$ which contradicts the last expression for any node $j \in \tilde{T}$. Thus, $\tilde{T} \supsetneq T$. Therefore, it must be that $\tilde{T} = T$, or that $(\tilde{S}_D, \tilde{a}) = (S_D^*, a^*)$. □

Proof of Lemma 1. Under δ^e we know that $\phi(\delta_e, i; G) = k > \phi(\delta_e, j; G)$ for all $i \in T$ and $j \notin T$. Similarly, under $\tilde{\delta}_e$, the allocation across the nodes must be proportional to their disruption value and such that $\phi(\tilde{\delta}^e, i; G)$ are the smallest possible for each node i. In particular, suppose $\tilde{\delta}_e$ such that $\phi(\tilde{\delta}_e, i; G) = z > \phi(\tilde{\delta}_e, j; G)$ for all $i \in N' \subsetneq N$ and $j \in N \setminus N'$ with z a positive real number. We are going to show that it must be that $N' = T$. Suppose it is not true. Then, it is easy to see that $N' \neq T \Rightarrow N' \subsetneq T \vee N' \supsetneq T$. Suppose $N' \subsetneq T$. This means that $\exists i \in T \setminus N' : D_i = 0$ and $\phi(\tilde{\delta}_e, i; G) < \phi(\tilde{\delta}_e, j; G)$ with $j \in N'$ such that $D_j > 0$. This contradicts δ_e since it required that $\phi(\delta_e, i; G) > \phi(\delta_e, j; G)$ for any $i, j \in T$ such that $D_i = 0$ and $D_j > 0$. The same argument is valid for the case $N' \supsetneq T$. Thus, $N' = T$ and therefore $\delta_e = \tilde{\delta}_e$. □

Proof of Lemma 2. It is trivial to see that there exists a player $q \in N$ such that $|\mathbb{P}'_{sq}| > |\mathbb{P}_{sq}|$, with $\mathbb{P}'_{sq}$ and $\mathbb{P}_{sq}$ respectively the sets of directed paths from s to q under G' and G and in particular $\mathbb{P}'_{sq} \supsetneq \mathbb{P}_{sq}$. Suppose $M_{sq}(G) \neq \varnothing$. We can show that $|M_{sq}(G)| \geq |M_{sq}(G')|$, or there is no player h such that $h \notin M_{sq}(G) \wedge h \in M_{sq}(G')$. Suppose by contradiction that there exists such player h. This means that $h \in \mathbb{P}'_{sq} \wedge h \notin \mathbb{P}_{sq}$, which is not possible since $\mathbb{P}_{sq} \subsetneq \mathbb{P}'_{sq}$. Therefore, there is no such player h. Finally, since for any node $k \in N$ such that $|\mathbb{P}'_{sk}| = |\mathbb{P}_{sk}|$ it must be that $M_{sk}(G) = M_{sk}(G')$, we obtain that $|M(G)| \geq |M(G')|$.

To see why a greater number of potential targets decreases $\alpha_i(D_i, \beta)$ for all $i \in T$, suppose a graph G and an equilibrium profile such that $|T| = m > 1$. We know that $\phi(S_D^*, i; G) = z > 0$ for all $i \in T$. Suppose we activate a new link obtaining $G'(N, L')$ and $T' \subseteq T$. If $T' = T$, the probability $\alpha_i(D_i', \beta) = \alpha_i(D_i, \beta)$ with $D_i' \in S_D^{*'}$ and for all $i \in T'$. Suppose $T' \neq T$ and such that $|T'| < m$. We know that $T \supsetneq T'$ and that $D_i' > D_i$ by Proposition 1, thus $\alpha_i(D_i, \beta) < \alpha_i(D_i'\beta)$. $\square$

A.2 Star Vs Core-Periphery

Suppose n players in a star $G(N, L)$ with one producer and $n-1$ non-producers players. The producer produces one unit of good. In particular, the node $O(G) = \{i\}$ and i is also the center or hub of the star network, Moreover L is composed only by links ij with $j \neq i$. It is easy to check that the equilibrium profile is such that the unique target of A is i, i.e. $T = \{i\}$ is singleton or i is attacked by A with certainty. The expected welfare when i is the unique target of A is simply

$$W = \alpha(n, \beta)nf(1)$$

The following trade-off is intuitive. A structure with a unique potential target concentrates all the defence resources on a unique player making him particularly resilient to attacks. On the other hand, a network such as a star where all the flows depend on the existence of the hub leads to a maximal disruption, in the case of a successful attack.

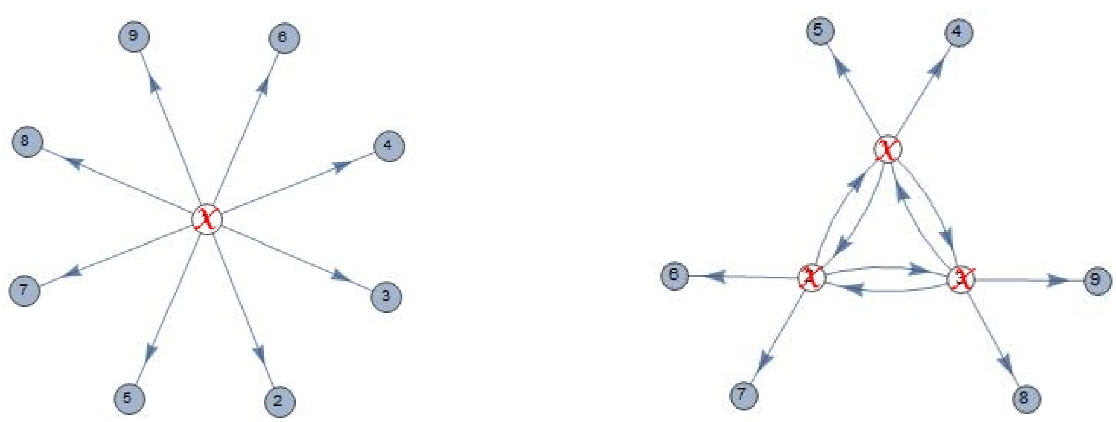

Figure 3: A star-shaped structure (left) and a core-periphery one (right). The producers are white colored while the peripheral non-producers are black. In the star $T = \{1\}$ and $D_1^* = 9$ while in the core-periphery structure $T = \{1, 2, 3\}$ and $D_i^* = 3$ for all $i \in T$.

We can now compare W to the welfare level when we assume a core-periphery graph where the same unit of good is produced in equal share by $1 < m < n$ producers. In particular, we define with core-periphery structure a directed network where the core is completely connected (there exists a bilateral link for each pair of nodes in the core), and $n - m$ nodes receive only one link from one of the nodes in the core (peripheral nodes). We assume that the $n - m$ non-producers are equally connected to one of the producers so that each producer is directly connected to $(n - m)/m$ non-producers. This structure makes each producer $j \in O(G)$ a potential target of A and in equilibrium we expect $D_j^* = n/m$ for each j. Intuitively, each producer disruption value is smaller than the previous case. However, the total defence resources are now spread across all the m producers instead of being concentrated in one and thus each individual target's probability of survival is smaller than in a star structure. The expected welfare is

$$W' = \alpha(n/m, \beta)nf(1) + (1 - \alpha(n/m, \beta))(m - 1)f\left(\frac{m-1}{m}\right)\left(\frac{n}{m}\right)$$

Therefore, $W \geq W'$ if and only if

$$nf(1)(\alpha - \alpha') - (1 - \alpha')(m - 1)f\left(\frac{m-1}{m}\right)\left(\frac{n}{m}\right) \geq 0 \tag{4}$$

where $\alpha \equiv \alpha(n, \beta)$ and $\alpha' \equiv \alpha(n/m, \beta)$. It is easy to check that the function on the left hand side of (4) is strictly concave with respect to β given n and m such that $n > m$, and positive for relatively low β values. In words, the star structure is preferred to the core-periphery one only for $\beta \leq \beta^*$ values. The threshold level $\beta^*(n)$ increases with n for a given number of producers m: By increasing n, the hub producer in the star graph increases her defence capability and therefore a star structure is beneficial despite the certain attack on the unique producer.

This finally suggests that under this setting a complete network with a directed path between any pair of nodes is the structure maximizing the welfare for any $\beta > 0$ level.[12]

[12]Loosely speaking, we can always obtain a complete network with such characteristics by increasing the core

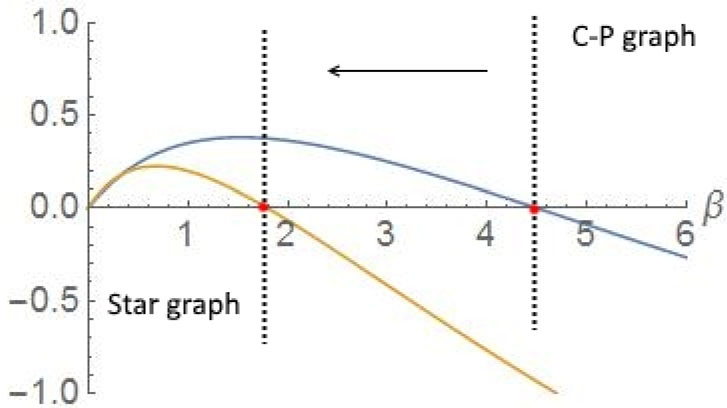

Figure 4: $\Delta W(\beta; n, m)$ function against β, assuming $\gamma = 1$ and linear $f(\cdot)$. For the blue function we assume $n = 9$ and $m = 3$, while for the yellow one $n = 9$ and $m = 6$.

In the plot of Figure 4, we compare the cases $n = 9$ and $m = 3$ (blue function), with $n = 9$ and $m = 6$ (yellow function). For relatively low attack intensity β, the star structure is favorite. However, all thing being equal, increasing n, the intersection happens for relatively higher β levels.

with respect to the peripheral to the limit where $m = n$, or each node is a producer of $1/n$ fraction of good and is connected to the rest of the players.

References

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, Dziubinski, Marcin, & Goyal, Sanjeev. 2015. Contagion risk and network design.

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.

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

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.

Major, John A. 2002. Advanced techniques for modeling terrorism risk. *The Journal of Risk Finance*, **4**(1), 15–24.

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.

Sims, Owen, & Gilles, Robert P. 2014. Critical Nodes In Directed Networks. *arXiv preprint arXiv:1401.0655*.

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