Decentralized Defence of a (Directed) Network Structure

Pelliccia, Marco

Defence and Peace Economics

DOI:

[10.1080/10242694.2019.1587679](https://doi.org/10.1080/10242694.2019.1587679)

Published: 17/08/2020

Peer reviewed version

[Cyswllt i'r cyhoeddiad / Link to publication](#)

Dyfyniad o'r fersiwn a gyhoeddwyd / Citation for published version (APA):

Pelliccia, M. (2020). Decentralized Defence of a (Directed) Network Structure. *Defence and Peace Economics*, 31(6), 659-676. <https://doi.org/10.1080/10242694.2019.1587679>

Hawliau Cyffredinol / General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Decentralised Defence of a (Directed) Network Structure

Marco Pelliccia*

Bangor Business School, Bangor University

Abstract

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

Keywords: Networks; Network defence, Security.

JEL classification: C72;

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

1 Introduction

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

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

Along the same lines, we propose a model of conflicts where a set of players (defendants) is connected by a directed network structure, and a (unique) attacker aims to maximally disrupt the network by attacking one of its nodes/players. Each defendant benefits from being part of the network as it gives him the possibility to receive goods produced by one or more peers. Each defendant is also endowed 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 de-

¹See 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).

²See Jackson *et al.* (2008).

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

We first show that when the attacker is strategic, nodes would share defence resources proportionally to their *criticality*³. More interestingly, we can show that under certain conditions the decentralized defence allocation is *efficient*; it coincides with the defence allocation which minimizes the expected network disruption. On the other hand, when the attack is probabilistic, this will not necessarily be the case; nodes which are more critical may end up being less defended than less critical ones. These results complete and, to some extent, challenge the existing literature.

The intuition behind the first result goes as follows. The directed nature of the network creates a topological ordering on each path connecting a player to a producer. This implies that for each player, his survival and that of any other player in the same path who is crucial to connect him to a producer are equally important. On the other hand, all things being equal, a strategic attacker would prefer to eliminate the most critical nodes. Under certain conditions, this will imply that (i) more critical nodes will receive relatively more defence resources from other peers (Proposition 1), and (ii) the interests of players in the same path will be coordinated, thereby aligning the decentralized allocation to a centralized one (Proposition 2). Loosely speaking, this is possible since the criticality of a player implies that a relatively large number of peers would be willing to defend him, thus he might be able to receive enough defence resources such that other less critical players would eventually become equally attractive to the attacker. When this does occur, this coincides with the allocation criteria of a planner. On the other hand, when the attacker is not strategic but probabilistic, each player would allocate defence to a peer only if the peer is essential to them, and proportionally to their probability of being attacked. In other words, the probability of being attacked, and not the players’ criticality, is essentially the unique element which affects a player’s defence allocation choice. Conversely, a planner would still take into account the criticality of each node. This is the reason why, in general, under random attack, decentralized and centralized defence allocations are likely to not coincide (Proposition 4).

Dziubiński & Goyal (2017), Acemoglu *et al.* (2016), Goyal & Vigier (2014), and Cerdeiro *et al.* (2017) are among the closest papers to ours. They study a sequential game in which a designer moves first and chooses a defence allocation, and in a second stage the *adversary* chooses how to allocate attack-resources across the nodes. In Cerdeiro *et al.* (2017), the authors also discuss how the designer could optimally design the network in order to solve

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

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

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

⁵There 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).

individual nodes might also inefficiently allocate defence resources. The first type has been the focus of the works mentioned above. Here, we concentrate on the second type of inefficiency, and in order to do so, we separate the individual defence-production problem from the resource-sharing one.

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

2 Network notation and definitions

A directed network $G(N, L)$ is composed by a set of nodes $N = \{1, \dots, n\}$ 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 $i_1, i_2, \dots, i_k$ such that $i_1 = i$ and $i_k = j$, and $i_1i_2, i_2i_3, \dots, i_{k-1}i_k \in L$. Two nodes are connected if there exists a path between them. A *cycle* is a path P_{ij} where $i = j$. We define the set of *predecessor nodes*, $B_i \subset N$, as the subset of nodes which can reach i by a path. Similarly the set of *successor nodes* or *follower*, $F_i \subset N$, the subset of nodes which can be reached from i through a path. We say that a node i is a *jq-middleman* node if and only if $i \in P_{jq}$ for any P_{jq} , or there are no paths from j to q which do not involve node i . Thus, we say that i is a middleman if and only if i is *jq-middleman* for at least one ordered pair (j, q) of nodes.⁶ We define a node i such that $F_i = \emptyset$, or who does not have any followers, a *sink* node. The *out-degree* of a node i , δ_i^+ , is the number of links departing from i , while the *in-degree*, δ_i^- , the number of links received by i . A *star graph* is a graph where a central node is connected to the rest of the players which are uniquely connected to him. A *core-periphery graph* is a graph similar to the star graph where a subset of players composes the core and are connected to the rest of peripheral players. A *directed acyclic graph* (or acyclic digraph) is a directed graph with no cycles. With abuse of notation, we indicate $G - i$ the graph obtained from G by removing the node i and any relative link. We finally define by $\mathcal{G}$ the set of directed networks.

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

3 Model

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

Definition 1. A directed network G with non-empty set of producer O is connected if and only if for each $s \in O$ there exists at least one path P_{si} to each $i \notin O$.

This definition will be particularly useful for some of the following results such as propositions 2 and 4. Each node is endowed of a unit of a divisible and transferable resource d which we call *defence resource*. We define $D_i = d_{ii} + \sum_{j \neq i} d_{ji}$, the total defence resources owned by i , where d_{ji} indicates the resource transferred by j to i . We assume that d is non-transferrable to third nodes, which is d_{ji} received by i from j cannot be transferred again to $q \neq j$.

3.1 Conflict

We analyse the two following scenarios:

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

We specify the technology of conflict. We assume that the attacker A always attacks with a constant intensity $\beta > 0$.

A node i owning D_i total defence, if attacked, survives with probability $\alpha(D_i)$ which is defined by a classic Tullock contest function⁸,

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

⁸See Tullock (2001).

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

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

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

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

where $B_i \cap O$ is the subset of producers who are also predecessors of i , or the producers who are connected to i . We assume $f(\cdot)$ concave such that $f'(\cdot) > 0$, $f(0) = 0$, and $f(nx_s)/f(x_s) \leq 2$ for any $n \geq 2$. In particular, the last condition guarantees that the marginal benefit of being connected to an additional producer is positive but small enough. This assumption simplifies the analysis and makes the problem tractable. If not holding, we could have two players who are equally crucial to other peers even if one connects few nodes to many producers while the other connects many nodes to few producers. This possibility complicates the analysis in a non trivial way and might be relevant only in cases where a single producer could not satisfy the demand of each receiver node. Summarising, player i benefits from being part of a component proportionally to the number of producers who can reach her *via* a path. If no such path exists, then there is no benefit from being part of the network. Thus, we can naturally compute the *network total value function* simply as $V(G) = \sum_{i \in N} v_i(G)$.

We remark that (1) is a generalization of the network value functions considered in Dziubiński & Goyal (2017), Goyal & Vigier (2014), and Cerdeiro *et al.* (2017). If we assume undirected graphs, or a node can reach any other node of the same a component, and each node is also a producer ($O = N$), then the argument of the function $f(\cdot)$ is essentially a multiple of the component's size.⁹ On the other hand, by assuming (1), we might also be able to describe cases where some path is not available, or where some node may be a simple receiver or an intermediary, i.e. where belonging to a network matters as long as it gives access to specific nodes by a path.¹⁰

⁹In such a case, we can also safely drop the condition $f(nx_s)/f(x_s) \leq 2$ for any $n \geq 2$ imposed on $f(\cdot)$.

¹⁰This may well describe the cases of trade networks, or infrastructure networks for example. Few countries

Finally, we define the node i 's *disruption value* $\tilde{V}_i = V(G) - V(G - i)$. In words, $\tilde{V}_i$ describes the potential impact of removing i from G on the defendants' valuation of the network. We can now describe the game in more detail.

3.2 Game setup

We consider a two-stage sequential game. In both S0 and S1, in the first stage the nodes simultaneously choose their defence allocation, while in the second stage one of the nodes is attacked. In S0 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 simultaneously chooses a strategy which is vector $x_i = (d_{i1}, \dots, d_{in})$ with $d_{ij} \geq 0$ and $d_{ii} + \sum_{j \neq i} d_{ij} \leq 1$. Thus the strategy space for each i is $S_i = [0, 1]^n$ and $S = S_1 \times \dots \times S_n$ the set of strategies. A defendant profile is $S_D = (x_1, x_2, \dots, x_n)$.

We focus now on the strategic scenario S1. Given S_D , the attacker chooses an attack profile $S_A = (\sigma_1, \dots, \sigma_n)$, where σ_i is the probability to attack node $i \in N$. When $\sigma_i = 1$, we refer to a pure strategy.

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

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

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

The expected payoff of attacker A under (S_D, S_A) is

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

All things being equal, $\phi(\cdot)$ has highest value when A attacks and destroys a node i such that $V(G - i) = 0$, and lowest when i is such that $\tilde{V}_i = v_i(G)$. In other words, the attacker's

own and export natural resources. The value of belonging to the trade network of a natural resource is linked exclusively to the existence of a trade path from the producer to the final country-consumer.

expected payoff increases with the chances of winning the conflict and with the expected disruption caused by the elimination of a target node.

A strategy profile (S_D^*, S_A^*) is a sub-game perfect Nash equilibria (SPNE) 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
- $\phi(G, S_D^*, S_A^*) \geq \phi(G, S_D^*, S_A)$ for all $S_A \neq S_A^*$.

We focus on the 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 , we expect 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 for any pair i and j attacked with positive probability,*

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

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

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

$$(1 - \alpha_i(D_i^*)) \tilde{V}_i = (1 - \alpha_j(D_j^*)) \tilde{V}_j$$

Thus, $D_i^* \geq D_j^*$, with equality holding only in the case $\tilde{V}_i = \tilde{V}_j$. Rearranging, we obtain the expression for D_i^* as stated above where it is easy to check that D_i^* is proportional to $\tilde{V}_i$. $\square$

The intuition is fairly simple. The attacker attacks more than one node with positive probability only if he finds them equally attractive. This means that if the attacker targets

two nodes of different disruption value with equal probability, it must be that the node with the higher disruption value is getting relatively more defence from other peers. Moreover, if multiple nodes are equally crucial to a player in order to be connected to a producer, he would find it optimal to transfer resources to them in a way that would make the attacker indifferent to attack either one of them; if the attacker was instead attacking one of these nodes with probability one and the player could divert some of his defence resources to the target, he would profitably do so.

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 $D = (D_1, \dots, D_n)$ with $D_i \geq 0$ for all $i \in N$ and $\sum_i D_i = n$, and, similar to the previous setting, the attacker chooses a distribution over the nodes in N given D . The expected payoff of the attacker is not changed while the CP's expected payoff is simply $\pi(D, S_A, G) = -\phi(D, S_A, G)$. 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. *Consider scenario S1. For any connected G where the nodes with highest disruption value are from the producer set, the centralized equilibrium profile (D^e, S_A^e) coincides with the decentralized one. Moreover, any decentralized equilibrium where the attacker randomizes over multiple nodes is unique.*

Proof: Recall that connected G means that each producer can reach any other non producer player in G . Define $m_q \geq 1$ the number of nodes depending on a node q to receive any good from a producer. Recall that, by the concavity assumption we made on $f(\cdot)$, $\tilde{V}_q > \tilde{V}_j$ implies $m_q > m_j$, or node q has strictly higher disruption value than node j only if q is crucial to more nodes than j in order to be connected to a producer.

The case where the CP allocates all the resources to a unique producer s and $\sigma_s^e = 1$ is easy to check. This would imply that

$$(1 - \alpha(n))\tilde{V}_s > \tilde{V}_i$$

for all $i \neq s$. Thus, being s crucial to any player in N , those would find it beneficial to send their own resource to s - any deviation from this would not change the response of A while

lowering the producer's defence. We are going to consider only cases when the centralized defence allocation implies that A would optimally randomize over at least two nodes, or best response $\sigma_s^e = \sigma_i^e > 0$ for at least another node $i \neq s$. Without loss of generality, consider the case of one producer s and one middleman node i . The argument for more than one producer and in general more targets would be similar. If in a decentralized setting the CP finds optimal to allocate resources over nodes such that $\sigma_s^e = \sigma_i^e = 1/2$, it must be that

$$(1 - \alpha(D_s^e))\tilde{V}_s = (1 - \alpha(D_i^e))\tilde{V}_i$$

Suppose $D^e \neq S_D^*$ and thus $D_i^e \neq D_i^*$. In particular, let's start by considering the case $D_i^e > D_i^*$. If this was the case, in the decentralized setting the attacker would optimally attack i with probability $\sigma_i^* = 1$. Moreover, since m_i players depend on i , they would all profitably send resources to i , thus $D_i^* = m_i$ and $D_i^e = m_i + \epsilon > D_i^*$, for some $\epsilon > 0$. This means that in the centralized setting it must hold

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

However, this is false for any $\epsilon \geq 0$ when $m_s > m_i$ and $m_s + m_i > n$, which is always satisfied. In particular, in order to be $\sigma_s^e = \sigma_i^e$, less than m_i resources needs to be allocated to i and more than $n - m_i$ to s . This is due to the fact that $\tilde{V}_q$ increases linearly with m_q while $\alpha(D)$ increases by less than a unit for each unit defence resource added. Thus, if $D_i^e = m_i + \epsilon$, optimal response of A in the centralized setting would be to attack s with probability one, a contradiction since S_A^e was already a best response. Consider now the case $D_i^e < D_i^*$. This implies that $\sigma_s^* = 1$. However, since G is connected and the producer s is assumed to be the node with highest disruption value, any follower of i would profitably divert resources to s , until $D_i^* = D_i^e$. Therefore, $D_i = D_i^e$, or $S_D^* = D^e$.

We finally show that when G is connected, there can only exist one equilibrium where A randomizes over more than one node. First, observe that when the attacker attacks multiple players with positive probability, there exists only one optimal defence allocation from their followers which makes the attacker indifferent to attacking any one of them. Thus, if there is more than one equilibrium in mixed strategy, it must be that there exists at least one defendant, say q , who, by deviating, would not strictly decrease his expected utility. Assuming G is connected and producers are the nodes with highest disruption value, the attacker would always attack one of the producers with positive probability. Since q is connected to the producers and benefits from receiving goods from these, any deviation would not be payoff neutral and would not be profitable. Therefore, under these conditions, an equilibrium

where the attacker attacks multiple nodes with positive probability must be unique. $\square$

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

Figure 1: Two examples where conditions for Proposition 2 do not hold and thus decentralized and centralized allocations might differ. In (a), player q is the node with highest disruption value despite not being a producer. This might lead to under-protection of q since the producers do not strictly benefit from defending q . In (b), producer j does not reach all the players. This might lead to under-protection of j . For instance, a planner might be able to allocate n resources making the producers equally attractive while this might not be reached in a decentralized setting.

More specifically, two conditions are sufficient in order to obtain the result. First, the producer(s) is(are) required to be the node(s) with highest disruption value. This is always the case if O is singleton but not necessarily when there are more than one producers. If this was not the case, we might have a middleman with highest disruption value who, in a decentralized setting, might not receive defence from his predecessor nodes thus potentially getting lower defence than in the planner's allocation. Again, we clarify this point with a simple example. Suppose two producers send goods to a unique middleman node q , who in turn connects them to other $n \geq 1$ players (see Figure 1a). It is easy to see that $\tilde{V}_q$, the disruption value of the middleman, is the highest in N . However, in a decentralized setting, we know that q would certainly receive defence only from himself and n followers since the producers do not strictly benefit from sending resources to q . This implies that if the CP would optimally allocate $D_s^e < 1$ to each producer and $D_q^e > n + 1$ to q , the centralized and decentralized defence allocations might differ.¹¹ By assuming networks where the producers exhibit the highest disruption value, we exclude these cases and guarantee the result.

Observe that we can easily quantify the highest inefficiency which could arise when the conditions stated in the Proposition do not hold. This is observed when a planner would allocate all n resources to one player q , while in a decentralized setting, q would receive resources only from $m_q < n$ peers. In such a case, we can calculate the difference in the expected disruption between the two settings as

$$\tau = \tilde{V}_q(\alpha(n) - \alpha(m_q))$$

where $n - m_q > 0$ if for instance $|B_q| \geq 1$, or q has at least one predecessor. For instance, when q is a middleman and the player with the highest disruption value, and a planner would allocate n resources to q , there is a decentralized equilibrium where all q 's predecessors and other players not depending on q do not transfer resources to q . In such a case, the *price of anarchy*, or the cost of a decentralized allocation in terms of higher expected disruption, is the highest. In the previous example, $n - m_q = |B_q|$, thus the highest price of anarchy would be $\tau = \tilde{V}_q(\alpha(n) - \alpha(n - |B_q|))$.

The second sufficient condition to obtain the result is that the producers need to be able to reach the rest of the nodes in N - the network is connected. Again, this is always the case when there is only one producer but not necessarily when more than one (see Figure 1b). If this condition is not holding, we might get an outcome similar to the case of two or

¹¹If $D_s^e = 0$ there exists a decentralized equilibrium where both producers send their resource to q since this would not affect their expected payoff. Therefore, the allocations certainly differ only if $D_s^e \in (0, 1)$.

more separate components of different size; a planner might still allocate defence such that the attacker is indifferent to attacking two or more players from distinct components while in a decentralized setting this might not happen when they are different in size. Intuitively, players from different components would not strictly benefit from sharing resources between them, thus decentralized and centralized allocation might differ.

Consider now scenario S0, and in particular assume that each node i can be attacked according to a probability distribution $P(i)$.

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 $s \in O$ and such that $\tilde{V}_j > 0$ proportionally to their probability of being attacked p_j . Moreover, if $p_i > 0$ for all $i \in N$ and G is a directed acyclic graph, the equilibrium defence allocation is unique.*

Proof: Suppose a path P_{sj} from a producer s to node j where each node in P_{sj} is essential to j to receive goods from a producer. Node j will allocate d_{jq}^* to $q \in P_{sj}$ in order to maximize the chances to receive the good. In particular, the probability to receive a good from s is β_1 , computed as

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

with $q \in P_{ij}$. Thus, d_{jq}^* must be proportional to p_q . Suppose a direct follower of j , node $k \in F_j$ such that $jk \in L$, who depends on j to receive the good from s . Node k will allocate resources to nodes in the path P_{sk} in order to maximize the probability β_2 , computed as

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

Thus, any amount $1 - d_{kk}^*$ optimally allocated by k to the nodes $q \in P_{ij}$ will maximize β_1 , thus d_{kq}^* will be proportional to p_q . On the other hand, if a player q is not essential to k despite being in a path P_{sk} - there exists at least another path P'_{sk} connecting k to a producer - then node k 's chances to receive the good from s do not depend on the existence of q , thus d_{kq}^* is

not proportional to p_q and it is likely to be equal to zero.¹² Finally, observe that if $p_i > 0$ for all $i \in N$ and G is directed and acyclic (DAG), then the equilibrium defence allocation is unique since in a DAG, each node depends on his predecessor nodes, if any, to receive goods from the producers. This implies that each node solves the problem of optimally allocating resources over a unique path as previously seen. Moreover, by assuming $\gamma \in (0, 1]$, the objective function has a unique maximizer since it would always consist of a sum of strictly concave functions. $\square$

In other words, players share defence resources with others in order to minimize the probability of disruption of paths connecting them to producers. Each node composing a unique path is equally essential to receive the good. Thus, the only element determining the defence received by a crucial node from other peers is the probability of being attacked, independently of his disruption value. The example in Figure 2 should clarify this point. Player 1 is the unique producer while player 2 is a middleman node. Consider beliefs $\{p_1 > p_2 > 0, p_3 = p_4 = 0\}$. Since player 1 is essential to all players and $p_1 > p_2$, we expect them to allocate more to 1 than 2 ($D_1^* > D_2^*$). Moreover, this is feasible since more players depend on 1 than 2, or $\tilde{V}_1 > \tilde{V}_2$, so there will be enough of them willing to satisfy the condition $D_1^* > D_2^*$. Consider instead beliefs $\{p_2 > p_1 > 0, p_3 = p_4 = 0\}$. We know that player 1 would always receive at least $d_{41}^* = 1$ and $d_{11}^* = 1$, so $D_1^* \geq 2$. Player 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 levels D_i^* not proportional to i 's disruption values.

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

This example anticipates the following result. Consider again a central planner who aims to minimize the expected network disruption. Which is, CP chooses the allocation $D^e = (D_1^e, \dots, D_n^e)$ which solves

¹²It is enough to assume $p_i > 0$ for at least one node i essential to k to be connected to any producer to guarantee $d_{kq}^* = 0$. If there is no such node, sending resources to q would never affect the payoff of k , thus we cannot exclude an equilibrium where $d_{kq}^* > 0$.

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

Proposition 4. *In general, for probability distributions $P(i)$ such that $p_i > 0$ for all $i \in N$, the CP's optimal allocation differs from the decentralized one.*

Proof: We prove the statement for a simple network of $n = 2$ nodes but the argument can be easily extended to the generic case of $n > 2$ nodes. Consider a producer s sending goods to a player i . Assume probability p_i of i being attacked and thus $(1 - p_i)$ the probability of s being attacked. If possible, the CP will allocate D_s and D_i such that the expected disruption $(1 - p_i)(1 - \alpha(D_s))\tilde{V}_s + p_i(1 - \alpha(D_i))\tilde{V}_i$ is minimized. Assume that the total resources available by CP are 2, thus $D_i = 2 - D_s$. Since $\tilde{V}_i = v_i$ and $\tilde{V}_s = 2v_i$, then D_s^e will satisfy

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

Consider now the decentralized equilibrium allocation where each node owns $d = 1$ defence resource. We know that the producer will always allocate a unit resource to himself, thus $d_{ss}^* = 1$, and we only need to check d_{is}^* . Node i will choose d_{is}^* in order to maximize the chances to receive and consume the good from s , or the probability $(1 - p_i)\alpha(1 + d_{is}) + p_i\alpha(1 - d_{is})$. Therefore, d_{is}^* will satisfy the condition

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

with equality holding if and only if $p_i < 1/2$, while for $p_i \geq 1/2$ we get the corner solution $d_{is}^* = 0$. When $p_i = 0$, it is trivial to see that $S_D^* = D^e$, or when the producer is attacked with probability one, the equilibrium and efficient allocation coincide. Let's assume a probability distribution over the nodes $P(i)$ such that $p_i > 0$ for both the nodes.

Observe that by increasing p_i from 0, the difference $D_s^e - D_s^* > 0$ increases, or the equilibrium allocation increasingly under-protects s compared to the efficient level. We show that this is true up to $p_i = 1/2$. Define $\tilde{p}$ the probability to get i attacked such that $D_s^e = 1$. This probability is unique and it is easy to check that $\tilde{p} > 1/2$. We also know that $D_i^* = 1$ for any $p_i \geq 1/2$ since the equilibrium would imply $d_{ii}^* = d_{ss}^* = 1$. This means that up to $p_i = \tilde{p}$, it must be $D_s^e > D_s^*$, while for $p_i \geq \tilde{p}$ it must be that $D_s^e \leq D_s^*$, with equality holding when

$p_i = \tilde{p}$. In other words, up to $p_i = 1/2$, the equilibrium allocation increasingly under-protects s compared to the efficient level. For $p_i > 1/2$, this difference is still positive but shrinking and eventually, when $p_i \geq \tilde{p}$, it becomes negative, i.e. the producer is over-protected in equilibrium compared to the efficient level. Therefore, the equilibrium allocation is efficient only if $p_i = \tilde{p}$ and then $D_s^e = D_s^* = 1$.

Finally, generalizing for the case of $n \geq 2$ nodes, there exists a unique distribution $\tilde{P}(i)$ where $p_i > 0$ for all $i \in N$ which guarantees $S_D^* = D^e$, and it is such that

$$\frac{\tilde{p}_i}{\tilde{p}_j} \frac{\tilde{V}_i}{\tilde{V}_j} = 1 \quad (2)$$

for all pairs $(i, j) \in N^2$. Such distribution will also imply an equilibrium efficient allocation where $D_i^* = 1$ for all $i \in N$ if $\tilde{V}_i \geq \tilde{V}_j$ for all i and j such that $j \in F_i$. In other words, if the probability of being attacked satisfies condition (2) for each pair of players and each player has weakly lower disruption value than his predecessors, if any, then they would not find it profitable to transfer their defence resources to other peers; by allocating resources to some predecessor who is attacked with lower probability, a player would not maximize the chance of receiving goods from a producer. Thus, $D_i^* = D_i^e = 1$ for all $i \in N$, or $S_D^* = D^e$. $\square$

Contrary to the strategic scenario, under S0 it is likely that the decentralized defence may result in a sub-optimal allocation. This is mainly due to the fact that the planner would take into account the disruption value of each node while individual players would base their allocations purely on $P(i)$. In particular, we can say that the equilibrium allocation is efficient only in two cases. First, when the attacker attacks the unique producer with probability one. In such case, it is intuitive to see that the nodes and the planner have all aligned objectives. Second, under a unique and specific probability distribution where the nodes are attacked with probability inversely proportional to their disruption values and these follow the topological order in G , which is a follower cannot have a disruption value strictly higher than his predecessor. In such case, we also know that the efficient equilibrium allocation will imply $D_i^* = 1$ for all $i \in N$. The intuition goes as follows. When nodes are attacked with probability inversely proportional to their disruption values, an individual's best response will be to allocate her own resource to herself. The planner will generally allocate resources proportionally to both the nodes' probability of being attacked and to their disruption values. Therefore, the (unique) probability distribution over the nodes which guarantees that the planner would assign equal "values" to each node is also the distribution where decentralized and centralized allocation coincide.

For example, consider again the network in Figure (2) and assume a random attack such that $p_i = 1/n$ for all i . Equilibrium allocation is $D_i^* = 1$ 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 $D_1^e > D_2^e > D_3^e = D_4^e$, thus $D^e \neq S_D^*$. On the other hand, suppose $p_1 = 1$. Then, $D_1^* = D_1^e = n$, or the decentralized and centralized equilibrium defence allocation trivially coincide. Consider now the probability distribution $\tilde{P}(i) = \{0.09, 0.18, 0.36, 0.36\}$. Then, $S_D^* = D^e$ and such that $D_i^* = 1$ for all $i \in N$. In fact $\tilde{P}(i)$ is the unique distribution where $p_i > 0$ for all i such that $S_D^* = D^e$.

5 Discussion

5.1 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 maximise 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.

Consider the following Utilitarian Welfare function given the equilibrium profile (S_D^*, S_A^*) ,

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

In other words, the welfare of players in N coincides with the expected total network value given targets T and defence allocation S_D^* . Note that any network G which in equilibrium maximizes $W(S_D^*, S_A^*, G)$, will also minimises $\phi(S_D^*, S_A^*, G)$. 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, 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 the attacker's expected payoff in equilibrium.

To see this more clearly, suppose that given a network G in equilibrium T was a singleton

set. According to Proposition 1, this implies that the target is a producer s and gets resources from all the rest of the nodes, thus $D_s^* = n$. 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, s will be part of the targets, and that $D_i^* < n$ for any player $i \in T'$ and thus for s too. Moreover, expected payoff of the attacker in G will be higher than in G' since $\phi(S_D^*, S_A^*, G') = (1 - \alpha(D_s^*))\tilde{V}_s > (1 - \alpha_s(D_s^*))\tilde{V}_s = \phi(S_D^*, S_A^*, G)$. 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 (see Figure 2).

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

We may also ask whether it is welfare improving to share the production among multiple nodes. Consider a star-graph G with one central producer node, say s , and $(n - 1)$ peripheral nodes. For n large enough, the unique equilibrium profile implies that the attacker will attack s with probability one and $D_s^* = n$. Consider now the following alternative architecture. Suppose that same unit produced by the unique producer in 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 $(n - 1)$ 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.

Let's assume that m is small enough such that (S_D^*, S_A^*) is such that A randomizes over the m producers. It is easy to see that when each core-node is connected to $n - 1$ nodes, we would always increase the expected welfare since $W(S_D^*, S_A^*, G) < W(S_D^*, S_A^*, G')$, where

G is the star-graph and G' the core-periphery graph with $m > 1$ producers in the core. In particular, this is true since

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

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

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

This is fairly intuitive since although in G' each producer has smaller defence ability than the unique producer in G , they also create smaller disruption in case of elimination by A , and the second effect always dominates the first.

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 a few core nodes and we make each of them the 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, 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 a too small m , 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 5: The core-periphery graph gives the highest welfare. The star graph is preferred to the not full core-periphery graph for relatively small m , or when the total production is shared among relatively few producers.

5.2 Variable costs

In the main section we made an important assumption regarding the cost of defence. Defendants owned a fixed amount of defence resources and their choices were not concerned with its production but only with its allocation. In other words, defence was a sunk cost. This might not be a safe assumption when describing security choices such as immunization decisions but a realistic one in other cases. For instance, in response to a specific threat, a government might not be able to produce new military resources in the short run but only be able to reallocate existing units to different “fields”.

In general, allowing for transfers between agents, we can identify two potential sources of inefficiency and thus divergences between decentralized and centralized security allocations.

Inefficiencies might arise due to differences between individual and centralized defence-production (over or under-investment in security by single players), and/or due to differences in the redistribution of existing defence resources across players. As previously noted, the first case has been extensively analysed in the literature. In the previous section we focused on the second source of inefficiency, and in order to fully disentangle the two, we assumed sunk costs and allowed for transfers between nodes.

For completeness, here we discuss the possibility of producing defence at a constant marginal cost. We start by analysing the case with no transfers of security between nodes. We assume hereafter that the attacker is strategic.

Consider a marginal cost $c > 0$ per unit of defence produced by a node. Without loss of generality, assume $v_i(G) = 1$ for all $i \in N$ receiving a good from a producer and a unique producer s . When not specified, the setting is the same as the one previously discussed. In the first stage, each defendant simultaneously chooses her own defence $d_i \in [0, \infty)$. In the second stage, the attacker chooses a target. In the next result, we show that an equilibrium exists only for high enough c and, when it does, it implies under-security compared to the efficient level.

Proposition 6. *For cost $c \geq \tilde{c}$, an equilibrium exists and implies that the attacker attacks the node(s) with highest disruption value. Moreover, the decentralized equilibrium defence profile exhibits under-protection compared to the centralized one.*

Proof: First, observe that there exists a unique level $d = \tilde{d}$ such that

$$(1 - \alpha(\tilde{d})) = \alpha'(\tilde{d})$$

since $(1 - \alpha(0)) = 1$ while $\lim_{d \rightarrow 0} \alpha'(d) = \infty$, both functions converge to zero for $d \rightarrow \infty$, and both decrease monotonically. Moreover, $\tilde{d}$ is finite since $(1 - \alpha(d))$ decreases at lower rate than $\alpha'(d)$. Define $\tilde{c}$ the level

$$\tilde{c} \equiv (1 - \alpha(\tilde{d})) = \alpha'(\tilde{d})$$

We show that for $c \geq \tilde{c}$ an equilibrium exists and it is such that $\sigma_i^* = 1$ with i the node with highest disruption value, d_i^* such that $\alpha'(d_i^*) = c$, and d_j^* for all $j \neq i$ such that $(1 - \alpha(d_j^*))\tilde{V}_j = (1 - \alpha(d_i^*))\tilde{V}_i - \epsilon$ with ϵ positive and infinitesimally small.¹³ It is easy to see that when $\sigma_i = 1$,

¹³If there exists more than one node i with highest disruption value, $\sigma_i = \sigma > 0$ for all these nodes i .

i would not find profitable to decrease d_i^* . Moreover, i does not find it profitable to increase her defence since when $c \geq \tilde{c}$, $d > d_i^*$ implies both $(1 - \alpha(d_i^*)) < c$ and $\alpha'(d) < c$ where $(1 - \alpha(d_i^*))$ would be the marginal benefit from attracting A to other locations. Thus, i does not profitably deviate from d_i^* . Any player $j \neq i$ does not find profitable neither to reduce d_j^* , attracting A toward her location which is also less protected, nor to increase production, since it would only increase cost without changing A 's target. Finally, given d_i^* and d_j^* for all $j \neq i$, the attacker clearly finds profitable to attack i with probability one, thus the defence and attack profile described define an equilibrium profile.

We now show that the equilibrium defence production is not efficient and in particular that in the decentralized setting the nodes under-invest in defence. It is easy to see that a central planner would choose the optimal defence level for the node i , d_i^e , such that $\alpha'(d_i^e) = c/\tilde{V}_i$, or taking into account the disruption value of node i . Similarly to the decentralized case, d_j^e for any $j \neq i$ will be the protection level satisfying $\tilde{V}_j(1 - \alpha(d_j^e)) = \tilde{V}_i(1 - \alpha(d_i^e))$ if possible, otherwise $d_j^e = 0$. The attacker's best response will be to randomize over j nodes and i . To see that this is an equilibrium, observe that any deviation where $d_i < d_i^e$ would clearly be not profitable. Moreover, since the planner would always try to make A indifferent over i and j , any allocation where $d_i > d_i^e$ would be dominated by D^e .

Finally, this implies that the decentralized defence of i is smaller than the efficient level since, $d_i^e = kd_i^*$ with $k > 1$ for all $n \geq 2$ defined as

$$k \equiv \frac{\alpha'[c/\tilde{V}_i]^{-1}}{\alpha'[c/v_i]^{-1}}$$

Moreover, this also implies that $d_j^e \geq d_j^*$, with equality holding only when $d_j^e = d_j^* = 0$. $\square$

The intuition is simple. There are two reasons why a player might want to increase production of her own defence. First, if targeted by the attacker, this would increase the chance his chance of survival. Thus, the player might want to produce defence until the marginal benefit from increasing this probability is worth the cost of an extra unit of defence. Second, defence choices are strategic complementary since by increasing his own defence, a player might attract the attacker toward other locations. Thus, a player would increase his own defence if the marginal benefit from increasing his probability of survival when he succeeds in attracting A to other locations, from $\alpha(d)$ to 1, is greater than the cost to produce an extra unit of defence. When the cost c is relatively high, this second marginal benefit is smaller than the first one. This implies that the optimal defence of a targeted node would be the level equalizing the marginal benefit from increasing probability of survival to the marginal cost c .

Moreover, the target will necessarily be the player with highest disruption value since the rest of the players are always able to attract A toward this node even with lower defence levels.

For completeness, we show that for $c < \tilde{c}$, we do not reach an equilibrium. First, observe that if $c < \tilde{c}$, then $(1 - \alpha(\tilde{d})) = c > \alpha'(\tilde{d})$. Say that i was producing d_i such that $\alpha'(d_i) = c$ and again d_j where $(1 - \alpha(d_j))\tilde{V}_j = (1 - \alpha(d_i))\tilde{V}_i - \epsilon$. Production cost is now low enough to make it profitable to i to increase d_i in order to attract A toward j , or $1 - \alpha(d_i + \epsilon) > c$ with ϵ positive and small. This is true up to $d_i = \tilde{d}_i$ such that $(1 - \alpha(\tilde{d}_i)) = c$. If $\tilde{d}_i$ is produced by i , then best response of j will be to produce $\tilde{d}_j$ such that $(1 - \alpha(\tilde{d}_j))\tilde{V}_j = (1 - \alpha(\tilde{d}_i))\tilde{V}_i - \epsilon$, and again A would still find it profitable to attack i with probability one. However, at this point i would profitably decrease d_i to the initial level satisfying $\alpha'(d_i) = c$ condition, since it would not affect the attacker's response while bringing to the optimal levels of protection, and j consequently would decrease their defence levels too at the initial levels d_j . Therefore, we obtain a cycle dynamic and fail to reach an equilibrium profile.

Finally, observe that the result does not change when we allow for transfers of defence resources between players; a node benefits from producing and transferring resources to another node until the marginal benefit from increasing probability of surviving an attack of that player is equal to the marginal production cost, thus $D_i^* = d_i^*$ for all $i \in N$. However, by giving the possibility of transferring resources, we will always obtain multiple equilibria; each level $D_i^* > 0$ can be obtained by the individual contribution of i and/or of any node depending on i to receive goods from a producer. Since it is always the case that the node with highest disruption is such that $\tilde{V}_i > v_i$, we always face multiple equilibria.

In summary, assuming a strategic attacker and a marginal cost c high enough to guarantee the existence of an equilibrium, we expect a decentralized production of defence which is strictly lower than the efficient one. This is the case irrespective of the possibility of sharing resources between nodes. This result, together with those in the previous section, suggest that in a decentralized and strategic setting, inefficiencies in security are likely to arise from individual production choices and not from the sharing of resources. Moreover, the inefficiency expected from an individual player is proportional to the player's decentralized equilibrium production. Which is $d_i^e - d_i^* \geq d_j^e - d_j^*$ for any pair of nodes such that $d_i^* \geq d_j^*$. This suggests that, all things being equal, a structure with a relatively large number of middleman nodes (e.g. a line network) would present larger discrepancies between decentralized and centralized total protection than a network with fewer middleman nodes.

6 Conclusion

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

We have studied a game from the same family where connected players are endowed of defence units which can be shared between them. We show that if the attacker is strategic (S1), the decentralized allocation of defence resources may be efficient, or it may coincide 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 the players in the network, in S0 they do not since the likelihood of an attack on a player is independent of his disruption value. These results lead us to the conclusion that under strategic scenarios network structures may coordinate individual defence choices to efficient allocations by imposing a common goal on the agents, i.e. survival of network flows.

We also discuss how the network architecture may impact the final welfare of the defendants. Reducing the number of middleman (non producer) players, or players which are crucial to the flow of the goods through the network, is always welfare improving. Core-periphery structures with producers as core players 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 player linked to other core players but only to a fraction of peripheral ones) may be optimal (second-best) only when the core is large enough and the cost per-link relatively small.

Finally, we explore the impact of a variable cost of defence production. In line with part of the literature, we show that when an equilibrium exists, it implies under-investment in defence by each player. This is a direct consequence of the fact that players fail to internalize the impact of their elimination on the rest of their peers. This result, together with the previous ones, suggest that under certain conditions inefficiencies in decentralized security choices arise at individual production level and not at the redistribution of existing defence resources.

References

- Acemoglu, Daron, Malekian, Azarakhsh, & Ozdaglar, Asu. 2016. Network security and contagion. *Journal of Economic Theory*, **166**, 536–585.
- Alpcan, Tansu, & Başar, Tamer. 2010. *Network security: A decision and game-theoretic approach*. Cambridge University Press.
- Aspnes, James, Chang, Kevin, & Yampolskiy, Aleksandr. 2006. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *Journal of Computer and System Sciences*, **72**(6), 1077–1093.
- Bier, Vicki. 2006. Game-theoretic and reliability methods in counterterrorism and security. *Pages 23–40 of: Statistical Methods in Counterterrorism*. Springer.
- Bier, Vicki, Oliveros, Santiago, & Samuelson, Larry. 2007. Choosing what to protect: Strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*, **9**(4), 563–587.
- Cerdeiro, Diego A, Dziubiński, Marcin, & Goyal, Sanjeev. 2017. Individual security, contagion, and network design. *Journal of Economic Theory*, **170**, 182–226.
- Dziubiński, Marcin, & Goyal, Sanjeev. 2017. How do you defend a network? *Theoretical Economics*, **12**(1), 331–376.
- Goyal, Sanjeev, & Vigier, Adrien. 2014. Attack, Defence, and Contagion in Networks. *The Review of Economic Studies*, **81**(4), 1518–1542.
- Heal, Geoffrey, & Kunreuther, Howard. 2004. *Interdependent security: A general model*. Tech. rept. National Bureau of Economic Research.
- Hong, Sunghoon. 2011. Strategic network interdiction.
- Israeli, Eitan, & Wood, R Kevin. 2002. Shortest-path network interdiction. *Networks*, **40**(2), 97–111.
- Jackson, Matthew O, et al. 2008. *Social and economic networks*. Vol. 3. Princeton university press Princeton.
- Kalai, Ehud, & Zemel, Eitan. 1982. Totally balanced games and games of flow. *Mathematics of Operations Research*, **7**(3), 476–478.

- Keohane, Nathaniel O, & Zeckhauser, Richard J. 2003. *The ecology of terror defense*. Springer.
- Kovenock, Dan, & Roberson, Brian. 2010. Conflicts with multiple battlefields.
- Kunreuther, Howard, & Heal, Geoffrey. 2003. Interdependent security. *Journal of risk and uncertainty*, **26**(2-3), 231–249.
- Lapan, Harvey E, & Sandler, Todd. 1993. Terrorism and signalling. *European Journal of Political Economy*, **9**(3), 383–397.
- Reijnierse, Hans, Maschler, Michael, Potters, Jos, & Tijs, Stef. 1996. Simple flow games. *Games and Economic Behavior*, **16**(2), 238–260.
- Sandler, Todd, & Enders, Walter. 2004. An economic perspective on transnational terrorism. *European Journal of Political Economy*, **20**(2), 301–316.
- Sandler, Todd, *et al.* 2003. Terrorism & game theory. *Simulation & Gaming*, **34**(3), 319–337.
- Smith, J Cole. 2010. Basic interdiction models. *Wiley Encyclopedia of Operations Research and Management Science*.
- Tullock, Gordon. 2001. Efficient rent seeking. *Pages 3–16 of: Efficient Rent-Seeking*. Springer.
- Varian, Hal. 2004. System reliability and free riding. *Pages 1–15 of: Economics of information security*. Springer.
- Washburn, Alan, & Wood, Kevin. 1995. Two-person zero-sum games for network interdiction. *Operations Research*, **43**(2), 243–251.
- Wood, R Kevin. 1993. Deterministic network interdiction. *Mathematical and Computer Modelling*, **17**(2), 1–18.
- Zhu, Shanjiang, & Levinson, David M. 2012. Disruptions to transportation networks: a review. *Pages 5–20 of: Network reliability in practice*. Springer.