UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Institute for New Economic Thinking

Cambridge-INET Institute

Working Paper Series No: 2015/04

CONTAGION RISK AND NETWORK DESIGN

Diego Cerdeiro (IMF)
Marcin Dziubiński (Warsaw University)
Sanjeev Goyal (University of Cambridge)

ABSTRACT

Individuals derive benefits from their connections, but these may, at the same time, transmit external threats. Individuals therefore invest in security to protect themselves. However, the incentives to invest in security depend on their network exposures. We study the problem of designing a network that provides the right individual incentives.

Motivated by cybersecurity, we first study the situation where the threat to the network comes from an intelligent adversary. We show that, by choosing the right topology, the designer can bound the welfare costs of decentralized protection. Both over-investment as well as under-investment can occur depending on the costs of security. At low costs, over-protection is important: this is addressed by disconnecting the network into two unequal components and sacrificing some nodes. At high costs, under-protection becomes salient: it is addressed by disconnecting the network into equal components.

Motivated by epidemiology, we then turn to the study of random attacks. The over-protection problem is no longer present, whereas under-protection problems is mitigated in a diametrically opposite way: namely, by creating dense networks that expose the individuals to the risk of contagion.

Contagion Risk and Network Design

Diego Cerdeiro*
Marcin Dziubiński†
Sanjeev Goyal‡

February 17, 2015

Abstract

Individuals derive benefits from their connections, but these may, at the same time, transmit external threats. Individuals therefore invest in security to protect themselves. However, the incentives to invest in security depend on their network exposures. We study the problem of designing a network that provides the right individual incentives.

Motivated by cybersecurity, we first study the situation where the threat to the network comes from an intelligent adversary. We show that, by choosing the right topology, the designer can bound the welfare costs of decentralized protection. Both over-investment as well as under-investment can occur depending on the costs of security. At low costs, over-protection is important: this is addressed by disconnecting the network into two unequal components and sacrificing some nodes. At high costs, under-protection becomes salient: it is addressed by disconnecting the network into equal components.

Motivated by epidemiology, we then turn to the study of random attacks. The over-protection problem is no longer present, whereas under-protection problems is mitigated in a diametrically opposite way: namely, by creating dense networks that expose the individuals to the risk of contagion.

Keywords: cybersecurity, epidemics, security choice, externalities.

---

*The International Monetary Fund. E-mail: dac54@cam.ac.uk.
†Institute of Informatics, Warsaw University. E-mail: m.dziubinski@mimuw.edu.pl
‡Faculty of Economics & Christ's College, University of Cambridge. E-mail: sg472@cam.ac.uk.
This paper is based on a chapter in Diego Cerdeiro's doctoral thesis submitted to Cambridge University in June 2014; the chapter was titled "Individual Security and Network Design". Sanjeev Goyal and Diego Cerdeiro were supported by European Research Area Complexity-Net. Marcin Dziubinski acknowledges support from the Foundation for Polish Science and Sanjeev Goyal also acknowledges support from a Keynes Fellowship and the Cambridge-INET Institute.

1 Introduction

Individuals derive benefits from being connected to others. These connections may, at the same time, transmit external threats. The internet reflects this tension:[1] connectivity facilitates communication but is also used by hackers, hostile governments and firms and 'botnet' herders to spread viruses and worms which compromise user privacy and jeopardize the functioning of the entire system.[2] Individuals are aware of these dangers and invest in security software. The incentives to invest in protection depend on exposure in the network and will generally depart from what is collectively desirable.

In this paper, our goal is to examine how network design can mitigate inefficiencies in protection.

There are $(n + 2)$ 'players'. The designer first chooses the network over the n nodes. Given this network, each of the n nodes (simultaneously) chooses whether to protect or not; protection carries a fixed cost. Finally, the adversary chooses a node to attack. If the attacked node is protected, then all nodes survive the attack. If the attacked node is not protected, then this node and all nodes with a path to the attacked node through unprotected nodes are eliminated. Nodes are assumed to derive benefits from their connectivity: the payoff of a node is increasing in the size of its surviving component. A node's net payoffs are equal to its connectivity payoffs less the amount spent on protection. The designer is utilitarian: he seeks to maximize the sum of nodes' payoffs. The adversary is intelligent, purposefully choosing the attacked node so as to minimize connectivity-related payoffs.

We start with a study of the first best design and defence profile. We show that for low protection costs, all nodes should be protected and any connected network is optimal. For intermediate costs of protection, the designer chooses a star network and protects its center only. The adversary then eliminates a single spoke of the star. If protection costs are high, the designer splits the network into equal size components and leaves all nodes unprotected. The adversary eliminates one of these components.

This sets the stage for the study of decentralized protection. We show that if defence is sufficiently expensive (so that no protection is first best), no protection is the unique equilibrium defence of any first best network. At the other extreme, if protection is sufficiently cheap (so that full protection is first best), there exist networks that implement the first best in every equilibrium. Departures from first best welfare will therefore arise

[1] In the United States, the Department of Homeland Security (DHS) is responsible for cybersecurity. Its mission statement reads, "Our daily life, economic vitality, and national security depend on a stable, safe, and resilient cyberspace. We rely on this vast array of networks to communicate and travel, power our homes, run our economy, and provide government services."

[2] Moore et al. (2009) estimate that in 2009, roughly 10 million computers were infected with malware designed to steal online credentials. The annual damages caused by malware are very large: in the US the annual costs of identity theft are estimated at 2.8 billion USD. These large costs have led to the emergence of a large software security sector. Intel bought McAfee in 2010, for 7.68 billion USD (bbc.co.uk; 19 August 2010).

only for intermediate costs of protection; that is, when a center protected star is optimal. The designer cannot attain first best payoffs in equilibrium, as the only equilibria on star networks are those where either all or no node protects. In our main result (Theorem 1), we show that the designer can bound the welfare costs of decentralization by choosing the right topology.

We then consider the optimal design problem in more detail. When a center protected star is first best but all nodes protect in equilibrium, protection decisions involve negative externalities and exhibit strategic complementarities. Nodes have incentives to protect and divert the adversary's attack to other parts of the network. How can the designer induce some nodes to be eliminated in equilibrium? We show that connected networks are suboptimal to address the over-protection problem. When a connected network has an equilibrium achieving higher welfare than full protection, there always exists a disconnected network that welfare-dominates it. Thus, if the designer is to avoid the over-protection problem, he must disconnect the network and sacrifice some nodes.

The analysis summarized so far assumes that individual coordinate on equilibria that achieve maximum equilibrium welfare. In general, however, some of these networks may feature multiple equilibria that achieve vastly different welfare levels. How can the designer tackle potential coordination problems? To illustrate the issue, suppose that the costs of protection are such that maximum equilibrium welfare is achieved via full protection on a connected network. The network where nodes are arranged on a cycle has a full protection equilibrium. However, if the cost of protection outweighs the benefits of surviving in isolation, there is another equilibrium on this network where no node protects and the adversary brings down the entire network. We provide a necessary and a sufficient condition for a network to induce full protection in any equilibrium. Such networks are *sparse* in the following sense: they must feature a node that can block the adversary's attack, thus saving a large part of the network.

Epidemics of diseases such as influenza, AIDS and tuberculosis, have enormous costs in terms of human suffering.[3] In the case of diseases, it is more natural to suppose that 'attack' on the social network is random. We show that in the first best scenario, optimal network structures do not change with the nature of the external threat if some level of protection is optimal. That is, the designer chooses either a connected network with all nodes protected (if security is sufficiently cheap), or a center-protected star (for intermediate values of protection costs). When protection is expensive, the optimal unprotected network depends on the value of connectivity. For very convex value functions, the designer may 'risk it' by creating a very large component, an option that is obviously suboptimal under intelligent threats.

[3]There are 3 to 5 million cases of acute influenza and between 250,000 and 500,000 deaths are attributed to this infection, annually. In 2012, over 8.5 million people were infected with tuberculosis and 1.3 million deaths were attributed to it. In the same year, 2.3 million new cases of AIDS were reported and over 1.5 million people died due to the disease; over 36 million people have died due to HIV/AIDS so far (WHO (2013, 2014a, 2014b)).

The solution to the design problem with decentralized security stands in sharp contrast with the case of intelligent attack. First, the over-protection problem is no longer present. Secondly, the under-protection problem may need to be addressed in a diametrically opposite way. Facing an intelligent adversary, security choices exhibit complementarities, and to avoid an equilibrium where nobody protects the designer must choose relatively sparse networks. Under random attack, security choices feature both strategic complementarities (due to the value of being connected to surviving individuals) and substitutes – a node will simply not protect unless it is sufficiently exposed to the risk of contagion. Since a node must be exposed to potential contagion in order to protect, the designer may need to choose dense networks to induce protection.

The contribution of the paper lies at the intersection of economics and computer science literature. For an early contribution in the study of decentralized defence, see Kunreuther and Heal (2003). Aspnes et al. (2006) studies security choices by nodes in a fixed network when nodes only care about their own survival, attack is random, and both protection as well as contagion are perfect. The focus is on computing the Nash equilibria of the game. They provide approximation algorithms for finding the equilibria. In a recent paper, Acemoglu et al. (2013) study the incentives for protection in a setting when both defence and contagion are imperfect.[4] The present paper provides, to the best of our knowledge, the first systematic study of the problem of optimal network design when the nodes invest to protect themselves against attacks.

Our paper is related to a recent literature on network design. Goyal and Vigier (2014) study the problem of security in a setting where security and network design are both chosen by a single player. The results in the present paper highlight the large effects of decentralized defence for optimal design. In Goyal and Vigier (2014) the optimal design is a star network and optimal allocation of resources is exclusively on the central node. By contrast, when individual nodes choose security, the optimal design has to address problems of too much as well as too little protection. This best way to tackle over-protection is by disconnecting the network and sacrificing some nodes. Potential under-protection problems are addressed by creating equal components. Finally, coordination problems in security are mitigated through the creation of 'sparse' networks that contain critical nodes.[5]

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 presents the first best solution. Section 4 presents our main result on the welfare costs of decentralization. Section 5 discusses optimal design. In Section 6 we consider the case of random attack. We conclude in Section 7. All proofs are in the Appendix.

[4] There is also a very active research programme in financial contagion, see e.g., Blume et al. (2013), Acemoglu et al. (2015), Cabrales et al. (2013), and Elliot et al. (2014)).

[5] Baccara and Bar-Isaac (2008) study the optimal cross-holding of incriminating information in a criminal organization, exploring the tradeoff between cooperation enforcement and potential detection by an external authority. However, no protection technology is available to agents; the choice of security is central to our study.

2 The model

Let $N = \{1, \ldots, n\}$, $n \geq 2$, be the finite set of *nodes*. A *link* is a two-element subset of N. A *network* G is a set of links, $G \subseteq \{ij : i, j \in N, i \neq j\}$, where ij is short for $\{i, j\}$. The set of all networks over the set of nodes N is denoted by $\mathcal{G}(N)$. Given a set of nodes $U \subseteq N$, $G[U] = \{ij \in G : i, j \in U\}$ is the subnetwork of G induced by U. Additionally, given a set of nodes $X \subseteq N$, $G - X = G[N \setminus X]$ is the network obtained from G by removing all nodes from X together with the adjacent links. A *path in G between nodes* $i, j \in N$ is a sequence of nodes $i_0, \ldots, i_m \in N$ such that $i_0 = i$, $i_m = j$, $m \geq 2$, and $i_{k-1}i_k \in G$ for all $k = 1, \ldots, m$. Node j is *reachable* from node i in G if $i = j$ or there is a path between them in G. We denote this fact by $i \overset{G}{\longleftrightarrow} j$. A *component* is a maximal set of nodes $C \subseteq N$ such that for all $i, j \in C$, $i \neq j$, we have $i \overset{G}{\longleftrightarrow} j$. The set of components of G is denoted by $\mathcal{C}(G)$. We will abuse the terminology and use the term 'component' to refer to the subnetwork $G[C]$ induced by a component C, as well. Given network G and node $i \in N$, $C_i(G)$ denotes the component $C \in \mathcal{C}(G)$ such that $i \in C$. Network G is *connected* if $|\mathcal{C}(G)| = 1$.

A *network value function* (NVF) is a function that reflects how good the network is in the given context. We consider the following network value function:

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

where $f : \mathbb{R} \to \mathbb{R}$ is strictly increasing, strictly convex, and $f(0) = 0$. This form of network value functions is in line with Metcalfe's law, where the value of a connected network over x nodes is equal to x^2, as well as Reed's law, where the value of a connected network over x nodes is of exponential order with respect to the number of nodes (e.g. $2^x - 1$). It reflects the idea that each node derives additional utility from every node it can reach in the network.

Players. There are $(n + 2)$ players: the designer (**D**), the n nodes, and the adversary (**A**).

The timing. There are three rounds of the game:

1. **D** chooses the network $G \in \mathcal{G}(N)$.

2. Nodes from N observe G and choose, simultaneously and independently, whether to protect (1) or not (0). This determines the set of protected nodes Δ.

3. **A** observes the protected network (G, Δ) and chooses node $x \in N$ to infect. The infection eliminates all the unprotected nodes reachable from x in $G - \Delta$. Thus the set of eliminated nodes is $E_x(G|\Delta) = \{i \in N : x \overset{G-\Delta}{\longleftrightarrow} i\}$, if $x \notin \Delta$, and $E_x(G|\Delta) = \varnothing$, otherwise. This leads to the *residual network* $G - E_x(G|\Delta)$.

Payoffs. Payoffs to the players are based on the value of the residual network and costs of defence. The *gross payoff* to node $i \in N$ in network G is equal to $f(|C_i(G)|)/|C_i(G)|$, i.e. each node gets the equal share of the value of its component. The net payoff of a node is equal to the gross payoff minus defence spending. The protection has costs $c \in \mathbb{R}_{++}$. A removed node gets payoff 0. Node i's payoff in network G with defended nodes Δ and attack x is then equal to:

$$U^i(G,\Delta,x) = \begin{cases} \frac{f(|C_i(G-E_x(G|\Delta))|)}{|C_i(G-E_x(G|\Delta))|} - c & \text{if } i \in \Delta \\ 0 & \text{if } i \in E_x(G|\Delta) \setminus \Delta \\ \frac{f(|C_i(G-E_x(G|\Delta))|)}{|C_i(G-E_x(G|\Delta))|} & \text{otherwise.} \end{cases} \tag{1}$$

The designer aims to maximize social welfare, i.e. the sum of nodes' utilities, which is equal to the value of the residual network minus total costs of defence. Formally, the designer's payoffs are equal to:

$$U^{\mathbf{D}}(G,\Delta,x) = W(G,\Delta,x) = \sum_{i\in V} U^i(G,\Delta,x) = \left(\sum_{C\in\mathcal{C}(G-E_x(G|\Delta))} f(|C|)\right) - |\Delta|c. \tag{2}$$

The adversary is intelligent and aims to minimize gross welfare, i.e. the sum of nodes' gross payoffs, equal to the value of the residual network:

$$U^{\mathbf{A}}(G,\Delta,x) = -\sum_{C\in\mathcal{C}(g-E_x(g|\Delta))} f(|C|). \tag{3}$$

The game. The model described above leads to game $\Gamma = \langle P, (\precsim_i)_{i\in P}, (S_i)_{i\in P}\rangle$. The set of players is $P = N \cup \{\mathbf{D}, \mathbf{A}\}$. The set of strategies of player $\mathbf{D}$ is $S_{\mathbf{D}} = \mathcal{G}(N)$. A strategy of each node i is a function $\delta_i : \mathcal{G}(N) \to \{0,1\}$ which, given network $G \in \mathcal{G}(N)$, provides the protection decision $\delta_i(G)$ of the node. The set of strategies of each node $i \in N$ is $S_i = 2^{\mathcal{G}(N)}$. A profile of individual strategies of the nodes determine, given a network G, the set of defended nodes $\Delta(G) = \{i \in N : \delta_i(G) = 1\}$. The set of strategies of player $\mathbf{A}$ is a function $x : \mathcal{G}(N) \times 2^N \to N$ which, given network $G \in \mathcal{G}(N)$ and set of protected nodes $\Delta \subseteq N$, provides the attacked node $x(G,\Delta)$. The set of strategies of player $\mathbf{A}$ is $S_{\mathbf{A}} = N^{\mathcal{G}(N)\times 2^N}$. A strategy profile is a tuple (G,Δ,x) with the strategy choices of each of the players.[6] The outcome of strategy profile (G,Δ,x) is $(G,\Delta(G),x(G,\Delta(G)))$. The preferences of players $\mathbf{D}$ and $\mathbf{A}$ are determined by their utilities from the outcomes of strategy profiles. In the case of nodes we make an additional tie breaking assumption that in the case of utilities being equal, each node prefers to stay uninfected.

Equilibrium. We are interested in subgame perfect equilibria of game Γ, called equilibria, for short. Throughout the paper we will also refer to the subgame ensuing after

[6] We will represent the strategies of the nodes with the function providing the set of defended nodes, for short.

network G is chosen. We will denote this subgame by $\Gamma(G)$. We will abuse the notation by using the same letters to denote the strategies in $\Gamma(G)$ and in Γ (we will indicate whenever it is not clear from the context which game is considered).

It is important to note that, for the problem to be well defined, for any network G the subgame $\Gamma(G)$ must have an equilibrium. This is established by the following lemma.

Lemma 1. *For any network $G \in \mathcal{G}(N)$, $\Gamma(G)$ has an equilibrium.*

An immediate corollary of Lemma 1 is that the game Γ has an equilibrium.

3 First best outcomes

We start the analysis by characterizing the optimal choice if the designer can choose the protection profile as well as the network. Before we state the proposition characterizing the first best, we define the following set. Given $n \in \mathbb{N}$, let

$$Q^*(n) = \arg\max_{q \in \{1,\ldots,n\}} (q-2) f\left(\left\lfloor \frac{n}{q-1} \right\rfloor\right) + f(n \bmod (q-1)). \tag{4}$$

As will be clear in the next result, for sufficiently high protection cost the first best involves no protection. For a given network value function, elements in the set $Q^*(n)$ are related to the maximum number of components in the optimal unprotected network.

Proposition 1. *Suppose the designer chooses protection and design. Then*

(1) if $c \leq \min\{c_1(n), c_2(n)\}$, the network is connected and all nodes are protected,

(2) if $c_1(n) < c \leq c_3(n)$, the network is a star and only the centre is protected,

(3) if $c > \max(c_2(n), c_3(n))$, the network is unprotected and has $q-1$ components of size $\lfloor \frac{n}{q-1} \rfloor$ and one component of size $n \bmod (q-1)$ (if $n \bmod (q-1) > 0$),

where

$$c_1(n) = \frac{f(n) - f(n-1)}{n-1}, \tag{5}$$

$$c_2(n) = \frac{f(n) - (q-2) f\left(\left\lfloor \frac{n}{q-1} \right\rfloor\right) - f(n \bmod (q-1))}{n}, \tag{6}$$

$$c_3(n) = f(n-1) - (q-2) f\left(\left\lfloor \frac{n}{q-1} \right\rfloor\right) - f(n \bmod (q-1)), \tag{7}$$

with $q \in Q^*(n)$.

The response of the adversary to each of these networks and defence profiles is as follows. The adversary attacks any node in case (1), eliminates a spoke in case (2), and targets a node in a component of size $\lfloor \frac{n}{q-1} \rfloor$ in case (3).

When defence is sufficiently cheap all nodes will be protected, and the maximum gross payoff of $f(n)$ will be achieved through any connected network. For intermediate values of c, protecting all nodes is too costly but the damage caused by the adversary can be brought to a minimum with a center-protected star. When the cost of protection is large, no node is protected and an undefended network is optimal.

Consider, for example, the case of Metcalfe's Law (i.e. $f(y) = y^2$) with $n = 30$ nodes. If $c \leq 2.03$, first best is achieved through a connected and fully protected network. If $2.03 < c \leq 616$, then a centre protected star is optimal. Finally, if $c > 616$, then the designer chooses a network consisting of two components of 15 nodes.[7]

4 The price of decentralization

What are the welfare implications of decentralized protection decisions? We will use two measures: the price of stability and the price of anarchy.

The price of stability is defined as the fraction of payoff to the designer in the first best over the maximal payoff to the designer that can be attained in equilibrium of Γ (for the given costs of protection c):

$$\mathrm{PoS}(n, c) = \frac{W(G^{\mathrm{fb}}, \Delta^{\mathrm{fb}}, x^{\mathrm{fb}})}{\max_{(G,\Delta,x)\in\mathcal{E}(c)} W(G, \Delta(G), x(G, \Delta(G)))}, \tag{8}$$

where $(G^{\mathrm{fb}}, \Delta^{\mathrm{fb}})$ is a first best network and defence profile and x^{fb} is a best response to it by the adversary. The price of anarchy is defined as the fraction of payoff to the designer in the first best over the minimal payoff to the designer that can be attained in equilibrium of Γ (for the given costs of protection c):

$$\mathrm{PoA}(n, c) = \frac{W(G^{\mathrm{fb}}, \Delta^{\mathrm{fb}}, x^{\mathrm{fb}})}{\min_{(G,\Delta,x)\in\mathcal{E}(c)} W(G, \Delta(G), x(G, \Delta(G)))}. \tag{9}$$

It is easy to see that these measures provide, respectively, lower and upper bounds on the welfare costs of decentralization.

The following additional quantity will be used in the analysis of decentralized equilibria:

$$c_0(n) = \frac{f(n-1)}{n-1}. \tag{10}$$

We start by noting that there is no cost of decentralization if protection is sufficiently expensive or sufficiently cheap.

[7] The optimal number of components of an undefended network depends on the convexity of the value function. For $f(y) = y^2$ and $n \notin \{9, 15\}$, it consists of two large equal-size components and (if n is odd) an isolated node. For $n = 9$, the network with three equal-size components is the unique optimal undefended network; for $n = 15$, there are two optimal undefended networks: three equal-size components, and two size-7 components and an isolated node. Formally, we have that under this network value function $Q^*(9) = \{4\}$, $Q^*(15) = \{3, 4\}$, and $Q^*(n) = \{3\}$ for any $n \notin \{9, 15\}$.

Lemma 2. *If $c \leq \min\{c_0(n), c_1(n), c_2(n)\}$ or $c > \max\{c_2(n), c_3(n)\}$, then there exists network G such that the designer attains first best payoffs in every equilibrium of $\Gamma(G)$.*

Suppose that $c > \max\{c_2(n), c_3(n)\}$, so that the first best consists of an unprotected network G with the largest components being of size $\left\lfloor \frac{n}{q-1} \right\rfloor$, where $q \in Q^*(n)$. If protection is costly enough for the first best to be an optimal undefended network G, then in such a network any potential gains from connectivity are outweighed by the cost of protection.[8] No protection is the unique equilibrium defence profile of $\Gamma(G)$. If $c \leq \min\{c_1(n), c_2(n)\}$, the first best consists of a connected network G with all nodes protected. If, in addition, $c \leq c_0(n)$, then while there exist connected networks with equilibria involving less than full protection, there always exist networks such that all nodes protect in every equilibrium.

Lemma 2 therefore establishes that departures from first best may arise for two different reasons. Firstly, if the cost of protection is such that $c_0(n) < c \leq \min\{c_1(n), c_2(n)\}$. In this case, first best welfare is attained through full protection, but *any* network has an equilibrium where no node protects. Such a situation may only arise if the network value function features exponential growth, since $c_0(n) < c_1(n)$ requires that the value of a network over n nodes is at least twice as large as the value of a network over $n-1$ nodes.

Secondly, there will be departures from first best if the latter consists of a centre-protected star. This requires that the network value function f and the network size n be such that $c_1(n) < c_3(n)$.[9] Then, for $c_1(n) < c \leq c_3(n)$ first best is a centre-protected star, but this cannot be attained in equilibrium, as the only equilibria on star networks are those where either all or no node protects.

Lemma 3. *Let G be a star network. In any equilibrium of $\Gamma(G)$, either all nodes protect or no node protects.*

For large n, decentralization of protection cannot be problematic if $\lim_{y\to\infty} \frac{f(y)-f(y-1)}{y-1} = +\infty$. In this case, for sufficiently large n, $c < \min\{c_0(n), c_1(n)\}$ and therefore the price of anarchy equals one.[10] If, on the other hand, $\lim_{y\to\infty} \frac{f(y)-f(y-1)}{y-1}$ is bounded, then the wedge between first best and decentralized welfare will not vanish as n increases. Our main finding is that, for any network value function, the ability to choose the network topology allows to bound the welfare costs of decentralization. This is summarized by the theorem below.

[8] Formally, we have that $c_3(n) > \frac{f(\lfloor \frac{n}{q-1} \rfloor)}{\lfloor \frac{n}{q-1} \rfloor}$.

[9] Note, for example, that $c_1(n) > c_3(n)$ for any n if $f(y) = y^y$.

[10] This follows from the fact that if $c_0(n)$ is bounded then $c_1(n)$ is bounded. The condition $\lim_{y\to\infty} \frac{f(y)-f(y-1)}{y-1} = +\infty$ is satisfied, e.g., for $f(y) = y^\alpha$ and $\alpha > 2$, or $f(y) = \alpha^y - 1$ and $\alpha > 1$. These functions are sufficiently convex so that, for any finite protection cost, there exist a sufficiently large n such that a connected fully protected network is first best *and* there exist networks with full protection in any equilibrium.

Theorem 1. *For cost of protection c and network size n:*

(1) If $c < \min(c_0(n), c_1(n), c_2(n))$ or $c > \max(c_2(n), c_3(n))$, then $\text{PoA}(n,c) = \text{PoS}(n,c) = 1$.

(2) Suppose that $c_1(n)$ is bounded and $\min(c_0(n), c_1(n), c_2(n)) < c < \max(c_2(n), c_3(n))$. Then:

(a) $\lim_{n\to\infty} \text{PoA}(n,c) = \lim_{n\to\infty} \text{PoS}(n,c) = 1$, if $\frac{f(n)}{n}$ is unbounded.

(b) $\lim_{n\to\infty} \text{PoA}(n,c), \lim_{n\to\infty} \text{PoS}(n,c) \leq \frac{p}{p-c}$, if $\lim_{n\to\infty} \frac{f(n)}{n} = p < +\infty$ with $p > c$.

(c) $\lim_{n\to\infty} \text{PoA}(n,c), \lim_{n\to\infty} \text{PoS}(n,c) \leq \frac{p}{f(1)}$, if $\lim_{n\to\infty} \frac{f(n)}{n} = p \geq c$.

To gain intuition for point 2 in Theorem 1, it is useful to consider examples. Suppose that $f(y) = y^2$. Note that $c_1(n)$ is bounded, $\lim_{n\to\infty} c_1(n) = 2$. Moreover, $\frac{f(n)}{n}$ is unbounded. Hence, this network value function corresponds to case 2a. Since $\frac{f(n)}{n}$ is unbounded, for any cost c there exists a network size n such that $\frac{f(n-1)}{n-1} \geq c$. In this case, the designer can enforce full protection by choosing the right topology. Moreover, since $f(n)$ grows faster than n, the welfare implications of over-protection by $(n-1)$ nodes are negligible compared to connectivity payoffs.

In cases 2b and 2c, $\frac{f(n)}{n}$ is bounded. The value of connections becomes approximately linear as the network size increases. Consider, for example, $f(y) = y - \ln(y+1)$, so that $\lim_{n\to\infty} \frac{f(n)}{n} = 1$. If $c < 1$, then for sufficiently large n the designer can choose a connected network where all nodes protect. The average payoff across nodes in the first best is $\frac{f(n-1)}{n} - \frac{c}{n} \to 1$, while in decentralized equilibrium with overprotection $\frac{f(n)}{n} - c \to 1 - c$. Hence the price of anarchy is bounded above by $\frac{1}{1-c}$. If $c \geq 1$, then the designer can choose the empty network, and the price of decentralization is at most $\frac{f(n-1)-c}{(n-1)f(1)} \to \frac{1}{f(1)}$.

5 Decentralized security and optimal design

Our main result (Theorem 1) states that the welfare implications of decentralization can be bounded by choosing the network topology. In this section we turn the attention to the optimal design problem.

When protection decisions are decentralized, inefficiencies stem from two distinct sources: pure-externality problems and coordination problems. To illustrate the inefficiencies associated with pure-externality problems, let $f(y) = y^2$. Suppose that the first best is a centre-protected star (i.e. $c_1(n) < c \leq c_3(n)$). If $c > \frac{(n-1)^2}{n-1}$ and no spoke protects, the centre of the star strictly prefers not to protect. This is the underprotection problem due to positive externalities. If $c < n-1$, there is a unique equilibrium where all nodes protect. This is the overprotection problem due to negative externalities.

The second source of inefficiencies are coordination problems. Let $f(y) = y^2$, and suppose that the first best is a connected and fully protected network (i.e. $c \leq \min\{c_1(n), c_2(n)\}$). Consider a clique, i.e. a network where there is a link between any pair of nodes. There are two possible equilibrium outcomes. One where all nodes protect (attaining the social optimum), and another one where no node protects. The latter is due to the fact that it is not profitable for a node to protect if no other node survives in the network. Protection in this setting has features of threshold public goods: it is only profitable for the nodes to protect if there are sufficiently many other nodes protecting in the network.

In this section we analyze how the designer can mitigate the decentralization problems by choosing the right network topology. Given that, depending on the network, the subgame $\Gamma(G)$ may feature multiple equilibria, we will consider two polar cases. For any network G, the equilibrium of the subgame $\Gamma(G)$ that will be selected will be either *welfare-maximising* or *welfare-minimising*. Formally, for a given network $G \in \mathcal{G}(N)$, let $\mathcal{E}(c|G)$ denote the set of all equilibria of $\Gamma(G)$ under costs of protection c. An equilibrium (Δ, x) is *welfare-maximising* if

$$(\Delta, x) \in \arg\max_{(\Delta', x') \in \mathcal{E}(c|G)} W(G, \Delta, x(\Delta)). \tag{11}$$

An equilibrium (Δ, x) is *welfare-minimising* if

$$(\Delta, x) \in \arg\min_{(\Delta', x') \in \mathcal{E}(c|G)} W(G, \Delta, x(\Delta)). \tag{12}$$

Let $\mathcal{E}(c)$ denote the set of equilibria of the game Γ. An equilibrium $(G, \Delta, x) \in \mathcal{E}(c)$ is welfare maximising if

$$(G, \Delta, x) \in \arg\max_{(G, \Delta', x') \in \mathcal{E}(c)} W(G, \Delta(G), x(G, \Delta)), \tag{13}$$

and welfare minimising if

$$(G, \Delta, x) \in \arg\min_{(G, \Delta', x') \in \mathcal{E}(c)} W(G, \Delta(G), x(G, \Delta)). \tag{14}$$

Consider potential differences between design under centralized and decentralized protection. Any discrepancy between first best design and design under welfare maximising equilibria will reflect a pure-externality problem. Differences between design under welfare maximising and welfare minimising equilibria reflect coordination problems.

5.1 Metcalfe's Law

In this section we present the characterization of optimal networks for the case of Metcalfe's Law, that is, when $f(y) = y^2$. This functional form can be motivated, for example, by assuming that each individual in a component of size y has a piece of information that has a value of 1 to everyone (including herself). Thus, every node in a surviving component of size y receives a gross payoff of y, and the designer's gross payoff from this component is equal to y^2.

We find that three classes of networks are optimal under Metcalfe's law in a welfare maximising equilibrium. When protection costs are low, the designer keeps the network connected, and in the welfare maximizing equilibrium all nodes protect. As protection costs increase, the designer needs to construct a network such that not all nodes protect. To do so, s/he finds it optimal to save on protection at the expense of connectivity.

In particular, by creating a relatively large star component, and a smaller component. In equilibrium, only the center of the star protects and the smaller component remains unprotected and is eliminated. For large protection costs, the designer chooses the optimal unprotected network, where in decentralized equilibrium no node protects.

The following two quantities correspond to the sizes of the star and unprotected component of the network with partial protection:

$$
\begin{aligned}
s(n) &= \left\lfloor (n+1) - \sqrt{2n} \right\rfloor, \\
u(n) &= \begin{cases} n - s(n) - 1 & \text{if } (n - s(n) - 1)^2 \geq 2s(n) - 1 \\ n - s(n) & \text{otherwise} \end{cases} .
\end{aligned}
$$

Proposition 2. *Assume $f(y) = y^2$ and $n \geq 20$. If (G, Δ, x) is a welfare maximising equilibrium, then*

(1) if $0 < c \leq \min\{c_D(n), c_U(n)\}$ or $s(n) < c \leq c_U(n)$, G is connected and all nodes protect,

(2) if $c_D(n) < c \leq s(n)$, G features a star of size $s(n)$ and a component of size $u(n)$, and only the hub of the star protects,

(3) if $c > \max\{c_U(n), s(n)\}$, G features two components of size $\lfloor n/2 \rfloor$ and no node protects,

where

$$
\begin{aligned}
c_U(n) &= \frac{n^2 - [\lfloor n/2 \rfloor^2 + n \bmod 2]}{n}, \\
c_D(n) &= \frac{n^2 - [s^2(n) + (n - s(n) - u(n))]}{n - 1}.
\end{aligned}
$$

The response of the adversary to each of these networks and defence profiles is as follows. The adversary attacks any node in (1), eliminates $u(n)$ nodes in case (2), and eliminates a component of size $\lfloor n/2 \rfloor$ in case (3).[11]

Figure 1 illustrates the result for $n = 30$. If $c \leq 2.03$, full protection in a connected network is first best, and this can be attained in equilibrium on any such network. If the cost of protection is between $c = 2.03$ and $c = 616$, a center-protected star is first best. However, only two equilibria are possible on a star network: full protection (if $c \leq 30$) and no protection (if $c > 29$). Since no network can induce protection for costs above 30, the interest is in what the designer should choose if $c \in (2.03, 30]$.

If $2.03 < c \leq 13$, over protection implies a departure from first best but the designer finds it optimal to keep the network connected and have all nodes protect in equilibrium.

[11] The condition $n \geq 20$ is sufficient for the adversary not to attack a partially protected component in equilibrium. Within the finite number of cases not covered (i.e. for $n < 20$), we could not find a network such that this takes place on the equilibrium path.

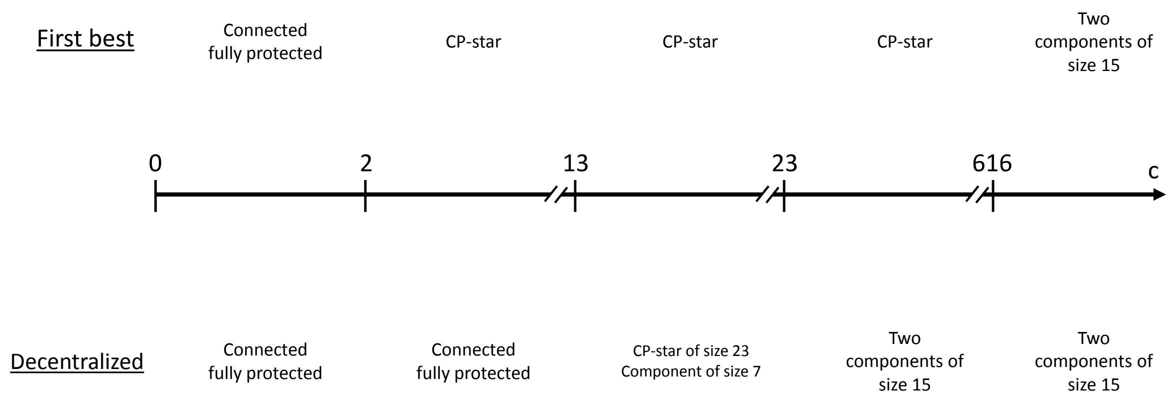

Figure 1: Optimal design as a function of protection cost c: $f(y) = y^2$ and $n = 30$.

For $c > 13$, the welfare costs of over-protection are severe enough to merit disconnecting the network to avoid it. The optimal network consists of a star of size 23 and a component of size 7. In equilibrium, only the center of the star protects and the adversary eliminates the component of size 7. While several nodes are compromised, many more save on protection.

For $c > 23$, however, the center of the star no longer finds it profitable to protect. Therefore, if $c \in (23, 30]$ the designer faces two options: either connecting all nodes and inducing full protection, or splitting the network and inducing no protection. Defence is sufficiently costly for the designer to prefer splitting the network into equal-size components and losing half of the nodes to the attack.

What networks are optimal if, for any network, a welfare minimizing equilibrium is chosen? Consider again the case with $n = 30$ nodes. For costs of protection above 23, the designer chooses the optimal unprotected network that consists of two components of size 15. Since $c > 23 > 15$, not to protect is a strictly dominant strategy for any node. Clearly, all equilibria on this network achieve the same level of welfare. Thus, in this case the set of optimal networks under welfare minimizing equilibria is the same as the one under welfare maximizing equilibria.

A similar argument extends to the case where $c \in [13, 23]$. That is, the optimal network under welfare maximizing equilibria attains the same welfare in any equilibrium. To see this, note first that a node in the small component of size $7 < c$ would never protect. Furthermore, if the center of the star does not protect, then the adversary would attack it, *even* if all spokes of the star protect. By eliminating the center of the star, the adversary causes a damage of at least $23^2 - 22 \cdot 1$.[12] This is vastly larger than the damage of 7^2 caused by attacking the unprotected component. Therefore, the center of the star protects in every equilibrium, and the adversary prefers to attack the unprotected component to eliminating an unprotected spoke.

Finally, if $c \leq 13$, the optimal network under welfare minimizing equilibria is connected and fully protected. In this range of costs, any connected network has a full protection equilibrium, but some of them have equilibria where not all nodes protect. However, for any $c \leq 13$ the designer can always choose a network that secures full protection in every equilibrium. This can be achieved, for example, by choosing the star network.

Formally, for a set of nodes N and a cost of protection c, let us denote with $\mathcal{G}^{full}(N, c)$ the set of connected networks such that all nodes protect in any equilibrium. That is,

$$\mathcal{G}^{full}(N, c) = \{G \in \mathcal{G}(N) : G \text{ is connected and } \Delta = N \text{ for any } (\Delta, x) \in \mathcal{E}(c|G)\}.$$

The following result establishes that this set of networks is not empty when in the welfare maximizing case the designer prefers full protection.[13]

Lemma 4. *Assume $f(y) = y^2$ and $n \geq 4$. If $c \leq c_U(n)$, then there exists a network where all nodes protect in every equilibrium.*

It follows from this result that, by choosing the right topology, in the welfare minimizing case the designer can avoid coordination problems and attain the same payoffs as in the welfare maximizing case. In terms of the price of decentralization, this means that the price of anarchy is equal to the price of stability. We summarize the welfare minimizing case in the following proposition.[14]

Proposition 3. *Assume $f(y) = y^2$ and $n \geq 4$. If (G, Δ, x) is a welfare minimizing equilibrium, then*

(1) if $0 < c \leq \min\{\hat{c}_D(n), c_U(n)\}$ or $\hat{s}(n) < c \leq c_U(n)$, G is in $\mathcal{G}^{full}(N, c)$ and all nodes protect,

(2) if $\hat{c}_D(n) < c \leq \hat{s}(n)$, G features a star of size $\hat{s}(n)$ and a component of size $u(n)$, and only the hub of the star protects,

(3) if $c > \max\{c_U(n), \hat{s}(n)\}$, G features two components of size $\lfloor n/2 \rfloor$ and no node protects,

[12] Formally, if y nodes protect in a star of size s, then the damage caused by eliminating the center equals $f(s) - yf(1)$. This is minimal for $y = s - 1$.

[13] If $n \in \{2, 3\}$ and $c \in (n - 1, c_U(n))$, the designer would like to induce full protection but any network has an equilibrium with no protection.

[14] In the next section we provide a necessary and a sufficient condition for a network to be in $\mathcal{G}^{full}(N, c)$ for general network value functions.

where

$$\begin{aligned} \hat{s}(n) &= \begin{cases} s(n) & \textit{if } 2s(n)-1 < u(n) \\ s(n)-1 & \textit{otherwise} \end{cases}, \\ \hat{c}_D(n) &= \frac{n^2 - [\hat{s}^2(n) + (n - \hat{s}(n) - u(n))]}{n-1}. \end{aligned}$$

Note that the size of the star, $\hat{s}(n)$, in the optimal network with partial protection may differ from the one the designer chooses under welfare maximising equilibria. This stems from the fact that, facing a given network and defence profile, different strategies of the adversary can be equilibrium strategies. Specifically, how the adversary decides, when indifferent between two attacks, can exacerbate the over-protection problem. For example, if $f(y) = y^2$ and there are $n = 32$ nodes, then the optimal partially protected network under welfare maximizing equilibria consists of a star of size 25 and a component of size 7. If only the center of the star protects, the adversary is indifferent between eliminating the small component (producing a gross welfare of $25^2 = 625$) and targeting a spoke of the star (yielding gross welfare of $24^2 + 7^2 = 625$). Two equilibrium outcomes are therefore possible: either all spokes protect, or no spoke protects (with the adversary attacking the small component in both cases). Clearly the former equilibrium is worse, and the designer responds to this by isolating a spoke of the star and thus reducing the size of the star to 24 nodes.

Figure 2 contrasts design under first best with design under decentralized protection for $f(y) = y^2$, as a function of the size of the network, n, and the protection cost, c. The parameter space (n, c) is partitioned into five regions. In region I, the first best is a connected network with all nodes protected. By choosing the right topology, the designer can attain first best payoffs. In regions II, III, and IV, the first best is a center-protected star. However, in a decentralized equilibrium either all nodes protect or no node protects in the star.

Facing this problem, the designer must choose to either keep the network connected, in which case all nodes must protect, or save on protection at the expense of connectivity. In region II, s/he opts for choosing a connected network. In turn, in region III a disconnected network with a center-protected star of size $\hat{s}(n)$ is optimal. It is important to note that, while this topology and protection profile can get the designer higher payoffs than both a fully protected connected network and an unprotected disconnected network, it may not be implementable in equilibrium. If $c > \hat{s}(n)$ the center of the star would not protect, and so the designer will choose either a connected network where all nodes protect (region II), or the optimal unprotected network that has two components of size $\lfloor n/2 \rfloor$ (region IV).

Finally, in region V the first best involves no protection, and this is implementable in equilibrium.

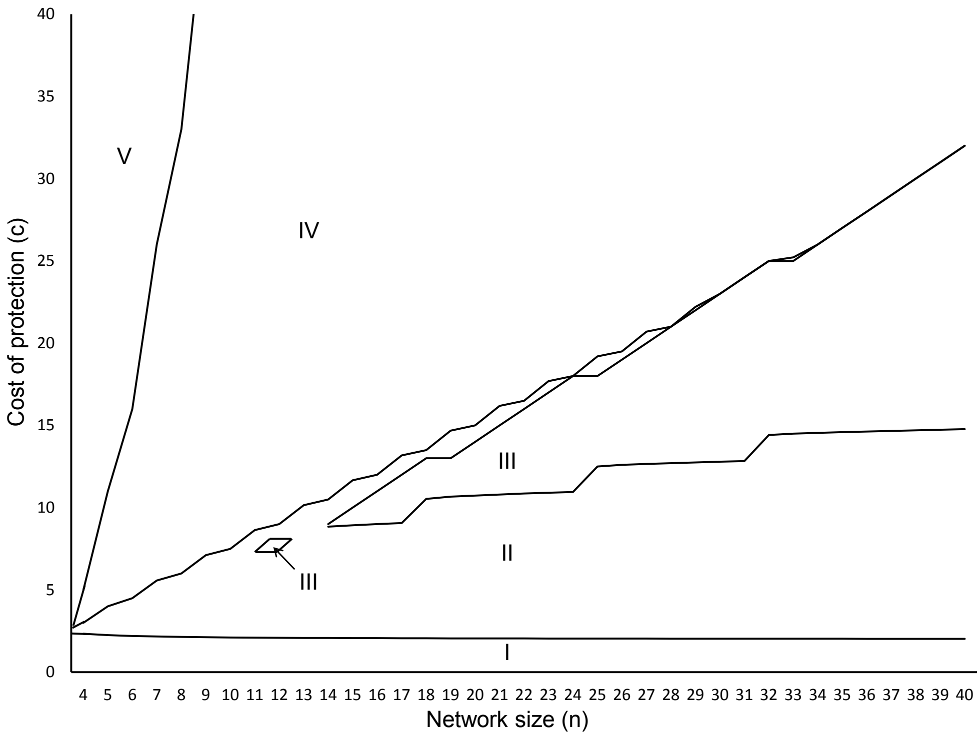

Figure 2: Optimal architecture if $f(y) = y^2$, as a function of network size n (horizontal axis) and protection cost c (vertical axis).

5.2 General network value function

The discussion of Section 4 established that departures from first best welfare arise if full protection is first best but is not implementable in equilibrium, and because the center protected star is not an equilibrium. If the first problem arises, its solution is simple: the designer will respond to it by choosing an optimal unprotected network. The second problem, on the other hand, is more challenging.

By Lemma 3, there are two equilibrium defence profiles on the star network. No protection is an equilibrium profile of the star if $c > f(n-1)/(n-1)$, whereas full protection is an equilibrium profile if $c \leq f(n)/n$. Note that for any $c > f(n)/n$, no network has protection in equilibrium, and thus the designer chooses in this case the optimal unprotected network. The non-trivial situation is therefore when $c \leq f(n)/n$ and in the best equilibria of the star all nodes protect. In the case of Metcalfe's Law explored above, we have seen that this is achieved by fragmenting the network and sacrificing a relatively small component. Here we show that this feature holds for any network value function, in the following sense. Specifically, we show that if the designer chooses a connected network, it must be that all nodes protect. The intuition is the following. Note that any connected network that has a partial protection equilibrium, also has a full protection equilibrium. If the benefits from connectivity are 'weak' enough for full protection to be worse than partial protection, then the designer can attain the same gross welfare (but with less protection) by luring the adversary with a relatively small component. Let us state the result and then provide an intuition for the formal proof.

Proposition 4. *Let G be an optimal network in a welfare maximising equilibrium of Γ. If some but not all nodes protect, then G is not connected.*

To illustrate the proof of the result, let us consider the following example. Suppose that there are $n = 20$ nodes, the cost of protection is $c = 0.8$, and the network value function is $f(y) = y - \ln(y + 1)$. The payoff to the designer under full protection is $f(20) - 20 \cdot c = 0.95$. This is clearly an equilibrium on any connected network, since $c < f(20)/20$. Suppose that the designer chooses a connected network where some nodes are eliminated in equilibrium. Consider, for example, the network depicted in panel (a) of Figure 3. The defence profile Δ^* is depicted such that protected nodes are surrounded by a square. Facing Δ^*, gross welfare is smaller if the adversary targets node i ($f(17) = 14.1$) than if s/he targets node k ($f(4) + f(15) = 14.6$). Note that if node i protected it would successfully divert the attack towards node k, since the gross welfare after attacking an unprotected neighbour of i is $f(18) = 15 > 14.6$. Consider then the strategy of the adversary that specifies attacking node i if nodes choose defence Δ^*, and attacking node k if nodes choose defence profile $\Delta^* \cup \{i\}$. Then Δ^* is an equilibrium defence. In particular, note that node i does not wish to protect, since in that case its payoffs are of $f(4)/4 - c < 0$.

This network thus avoids the over-protection problem, and achieves strictly higher welfare than a fully protected connected network ($f(17) - 2 \cdot c = 15.7 > 0.95$). Consider, however, re-designing the network into a star of 17 nodes and a cycle of 3 nodes, as shown in panel (b) of Figure 3. There is an equilibrium in this modified network where only the hub of the star protects and the adversary targets the unprotected cycle.[15] While gross welfare is $f(17)$ in both networks, the modified network features lower protection spending. This example points to the sub-optimality of partially protected connected networks. When they achieve higher welfare than the fully protected connected network, they are in turn dominated by a network which achieves the same gross welfare with lower spending on protection.

It follows from Proposition 4 that if the designer were to choose a network where some but not all nodes protect, then this network must be disconnected. This opens up two possibilities for the adversary's attack on the equilibrium path: either (i) s/he attacks an unprotected component, or (ii) s/he targets a partially protected component. While Case (ii) is the one that we have not been able to rule out for general network value functions, Case (i) is certainly a possibility, as shown above for the case of $f(y) = y^2$. It is important to note, however, that the option given by Case (i) of sacrificing a

[15] Facing this defence profile, the adversary prefers to attack the cycle (producing gross welfare of $f(17) = 14.1$) to eliminating a spoke of the star (gross welfare of $f(3) + f(16) = 14.8$). A node in the cycle could protect and thus divert the attack towards a spoke of the star. But this is not profitable, since $f(3)/3 - c < 0$.

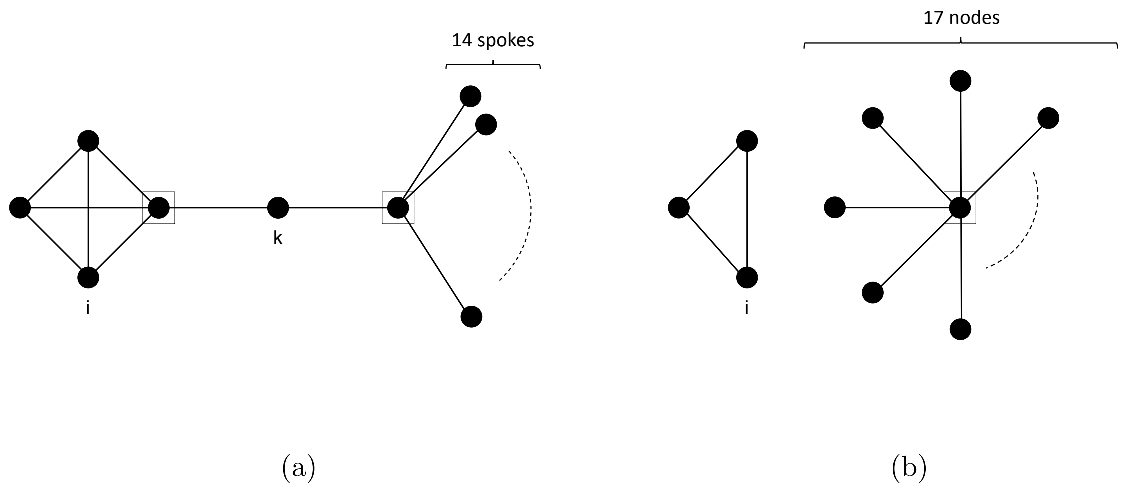

Figure 3: Addressing overprotection: $n = 20$ and $c = 0.8$, with $f(y) = y - \ln(y + 1)$. The network in panel (a) has an equilibrium where some but not all nodes protect. Panel (b) shows a disconnected network which achieves higher welfare.

relatively small component to save on protection in a larger component, which works under Metcalfe's Law, cannot work for some value functions.

To see this, consider any value function such that $f(y) > 2f(y - 1)$ (e.g. Reed's Law, $f(y) = 2^y - 1$). Let X be an unprotected attacked component, and Y an unattacked protected component, $|Y| > |X|$. As shown in the Appendix, if the network is disconnected no component is fully protected. Therefore, there is at least one node in Y that the adversary could eliminate. Since $|X| \leq |Y| - 1$, we have that $f(|X|) \leq f(|Y| - 1) < f(|Y|) - f(|Y| - 1)$, and so the adversary strictly prefers eliminating a single node of the largest component to eliminating all nodes in X, a contradiction. The intuition is simple: since $f(y) > 2f(y - 1)$, a single extra node in the largest component generates at least *twice* the value as the entire smaller component. Therefore, the designer will never be able to satisfy the 'appetite' of the adversary with a smaller component.

Let us next consider the problem of optimal design when, for every network, nodes and adversary coordinate on a welfare minimising equilibrium. How can the over-protection problem be addressed in these circumstances?

Observe first that, under welfare minimizing equilibria, if the optimal network is connected it must be that all nodes are protected. The intuition is as follows. Consider a connected network G which has a (welfare minimizing) equilibrium defence Δ where some but not all nodes protect. Even though the adversary is eliminating at least one node, the net payoff of nodes that protect is non-negative. It follows that network G must have another equilibrium defence where *all* nodes protect. Furthermore, by definition of Δ, the equilibrium with full protection cannot be worse. If Δ attains the same level of welfare as full protection, then the tie breaking assumption that nodes prefer to remain uninfected implies that full protection must be preferred by the designer as well. But then the designer can be better off by choosing a star network, as in any equilibrium of the star all nodes will protect.

The upshot of Proposition 4 therefore extends to welfare-minimizing equilibria (albeit for different reasons): the designer must disconnect the network if s/he is to avoid the over-protection problem. As stated above, if the network is disconnected then two things may happen on the equilibrium path. Either the adversary attacks an unprotected component (denoted as Case (i)), or s/he attacks a partially protected component (Case (ii)). Case (ii) can be ruled out if the network value function satisfies the following property.

Property 1. $f(1)y < f(y+1) - f(y)$ for any $y \geq 0$.

Property 1 is a condition on the convexity of the network value function. Functions that satisfy this property include, for example, Reed's Law ($f(y) = 2^y - 1$), and polynomial functions with exponent greater than or equal to 2 (i.e. $f(y) = y^\alpha$, $\alpha \geq 2$). We make the following remark.

Proposition 5. *Suppose f satisfies Property 1. Let (G, Δ, x) be a welfare minimizing equilibrium of Γ. If G is not connected, then the adversary attacks an unprotected component.*

The proof works by contradiction. Suppose that in a welfare minimizing equilibrium the network is disconnected but the adversary attacks a component where some nodes protect. Let us denote this component by P. Clearly, P cannot be fully protected. If P was fully protected and the adversary attacks it, then it must be that *all* components are fully protected, and the designer can be better off by choosing a connected star where all nodes protect in any equilibrium. Thus, P is not fully protected and the adversary eliminates at least one node of P.

Let us say that a component is 'large' if its size is such that $\frac{f(|C|)}{|C|} \geq c$, and that it is 'small' otherwise. Clearly, if a small component exists then eliminating it is always feasible for the adversary, since any node in a small component does not protect in any equilibrium. Moreover, there must exist one such small component if the optimal network is disconnected. Otherwise the defence profile where all nodes protect is an equilibrium defence, which, by definition, cannot be worse. But then the designer can attain strictly higher payoffs by choosing a connected star where all nodes protect in any equilibrium.

We next observe that if P has only one unit of protection, then the adversary must strictly prefer an attack on P to attacking a small component. To see this, note first that P must have the structure shown in Figure 4. The adversary eliminates the set X of nodes, and if a node in X protects then the adversary attacks a node in Y. For an eliminated node i who has a protected neighbour, not to protect is a best response only if the adversary disconnects protected nodes in P if i protects. Otherwise node i could get *at least* the same payoffs of protected nodes of P by protecting. Note, in particular, that component P has at least two units of protection and the adversary weakly prefers an attack on X or Y to an attack on any other component. Clearly, if P had only one unit of protection the adversary would *strictly* prefer an attack on P than an attack on any other component.

The proof is finalized with the following step. We show that it must be possible to construct another equilibrium defence, Δ', where in any 'large' component some but not all nodes protect, and the adversary attacks a 'small' component. By definition of Δ being a welfare minimizing equilibrium, the new equilibrium with defence Δ' cannot be worse. Moreover, by the observation of the previous paragraph, under Δ' there must be at least two protected nodes in component P. But then the designer can modify the original network by changing all components into stars. If Property 1 holds, then in any equilibrium on the modified network only the centres of the large components protect and the adversary eliminates a small component.[16] This attains the same gross welfare as Δ' does in the original network, but with strictly less protection spending. Hence the original network G cannot be optimal.

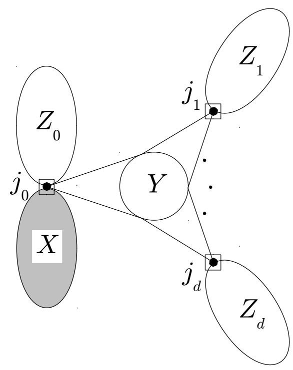

Figure 4: The structure of the partially protected and attacked component

As a corollary of Proposition 5, we have the characterization of the optimal networks under welfare minimizing equilibria for any network value function that satisfies Property 1. In equilibrium, if the network is connected then all nodes protect. If it is disconnected and some nodes protect, then it consists of large centrally protected stars and small unprotected components, with the adversary eliminating a small component. If it is disconnected and unprotected, it is the optimal unprotected network.

Consider, for example, the case of Reed's Law, $f(y) = 2^y - 1$. Since, as argued above, under Reed's Law the option of luring the adversary with a small component cannot work, optimal design under welfare minimizing equilibria is as follows.

Corollary 1. *Assume $f(y) = 2^y - 1$. Let (G, Δ, x) be a welfare minimizing equilibrium.*

(1) If $c \leq \frac{2^{n-1}-1}{n-1}$, G is connected and all nodes protect.

(2) If $c > \frac{2^{n-1}-1}{n-1}$, G features two components of size $\lfloor n/2 \rfloor$ and no node protects.

[16] If Property 1 holds, an equilibrium where the large components are periphery-protected stars may be possible.

Theorem 1 states that if $c < \min\{c_0(n), c_1(n), c_2(n)\}$, then the price of anarchy equals one. That is, first best involves full protection and there exist networks that attain full protection in every equilibrium. What is the structure of these networks?

Let us start by observing that if the cost of protection is very low, $c \leq f(1)$, then any node that is attacked is better off by protecting, regardless of the protection decisions of other nodes. Therefore, in any equilibrium outcome all nodes are protected. Since $f(1) < \min\{c_0(n), c_1(n), c_2(n)\}$, $c \leq f(1)$ implies that the first best is a connected and fully protected network. The designer can attain first best payoffs in decentralized equilibrium by choosing any connected network.

What if costs of protection are low (so that first best involves full protection) but not very low: $f(1) < c \leq \min\{c_0(n), c_1(n), c_2(n)\}$? As we discussed already, there are connected networks were inefficient equilibrium outcomes are possible, as nodes may fail to coordinate on the efficient equilibria (e.g. no protection on the cycle, when full protection is the first best). However, this problem can be solved by choosing the right topology of the network. Below we provide a necessary condition and a sufficient condition for a connected network to have full protection in any equilibrium outcome under costs of protection $c \leq \frac{f(n-1)}{n-1}$.[17]

Definition 1 (k-critical node)**.** Node $i \in N$ is *k-critical* in connected network G if the largest component in $G - \{i\}$ is of size k.

Loosely speaking, the importance of a node as a barrier against contagion due to an intelligent attack is decreasing in its criticality. For example, any node in a complete network is $(n-1)$-critical.[18] On the other hand, the centre of a star is 1-critical.

Proposition 6. *Consider a network G, and let k be such that $\frac{f(n-k)}{(n-k)} \geq c$.*

(1) If all nodes protect in every equilibrium of $\Gamma(G)$, then G has a k-critical node.

(2) If for all $i \in N$, i is k-critical or has a link to a k-critical node, then all nodes protect in every equilibrium of $\Gamma(G)$.

In essence, the presence of a k-critical node, with $\frac{f(n-k)}{n-k} \geq c$, rules out equilibrium outcomes where no node protects: each k-critical node has incentives to protect if no other node protects. However, it is not sufficient for having full defence in any equilibrium outcome. Consider the network depicted in Figure 5a. Let $f(y) = y^3$ and $c \in (81, 100]$. The largest component in $G - \{i\}$ is of size 9, and so i is 9-critical. Note that with $k = 9$, $\frac{f(n-k)}{n-k} = 100 \geq c$. Consider the defence profile shown in the figure. Facing this defence profile, the adversary generates a loss of $19^3 - (19-8)^3 = 5,528$ if node j is targeted,

[17]If $c > \frac{f(n-1)}{n-1}$, then every network has an equilibrium with no protection.

[18]In fact, any node in a d-connected network, $d \geq 2$, is $(n-1)$-critical. A network is d-connected if there is no set of $l < d$ nodes whose removal disconnects the network and the network can be disconnected by removing a set of d nodes (see e.g. Bollobás (1998)).

and a loss of $19^3 - (9^3 + 9^3) = 5,401$ if node i is attacked. Hence the best response of the adversary eliminates node j, which thus earns payoff 0. If j chooses to protect, then the adversary can generate a loss of only $19^3 - (19-7)^3 = 5,131$ if s/he attacks a node of the clique to which j belongs, and therefore prefers to attack node i when j protects. Thus, if it deviates to protection, payoffs of j are of $9^2 - c < 0$. Additionally, since each of the protected nodes earns positive payoff, none of them is better off by choosing no protection, as in any best response of the adversary they would be eliminated. Thus, the defence profile shown is indeed an equilibrium profile.

Note that a k-critical node cannot be eliminated in equilibrium, or otherwise it would profitably deviate by protecting. It follows from this observation that if G has a k-critical node and there exist unprotected nodes in an equilibrium of $\Gamma(G)$, then none of the nodes who are eliminated in equilibrium can have a link to a k-critical node. Thus, a sufficient condition for full protection to be the unique equilibrium outcome on a network is that every node is k-critical or has a link to a k-critical node. Figure 5b provides an example where the sufficient condition stated in the second part of Proposition 6 holds. Node i is 3-critical, and has links to all other nodes. Thus if $f(y) = y^3$ and $c \in (0, 16^2)$, any equilibrium of $\Gamma(G)$ has all nodes protected. This condition is not necessary, as illustrated in Figure 5c. When $f(y) = y^3$ and $c \in (0, 13^2)$, in any equilibrium outcome all nodes protect.

6 Random attack

To understand the effect of adversarial intelligence on the problem faced by the designer, in this section we consider the case where the identity of the node attacked is independent of its position in the network and protection status. In particular, the attack studied in this section is random in the following sense: a randomly picked node $i \in N$ is targeted. The payoffs of nodes and designer are modified in obvious ways to reflect expected utilities.[19]

[19]This model of random attack is the appropriate benchmark to study the effects of the adversary purposefully choosing one node to attack. An alternative model of random attack consists of assuming that every node fails independently with probability $1/n$. To see that the two specifications of random attack are different in a meaningful way, suppose that $f(y) = y^2$, $n = 4$, and $c > f(4) = 16$ so that investing in protection cannot be optimal for the designer. If every node fails independently with probability 1/4, then a connected network achieves the highest welfare, of $\left(1 - \frac{1}{4}\right)^4 \cdot f(4) = 5.0625$. Clearly, a connected network cannot be optimal if a randomly picked node is attacked, as it yields zero welfare.

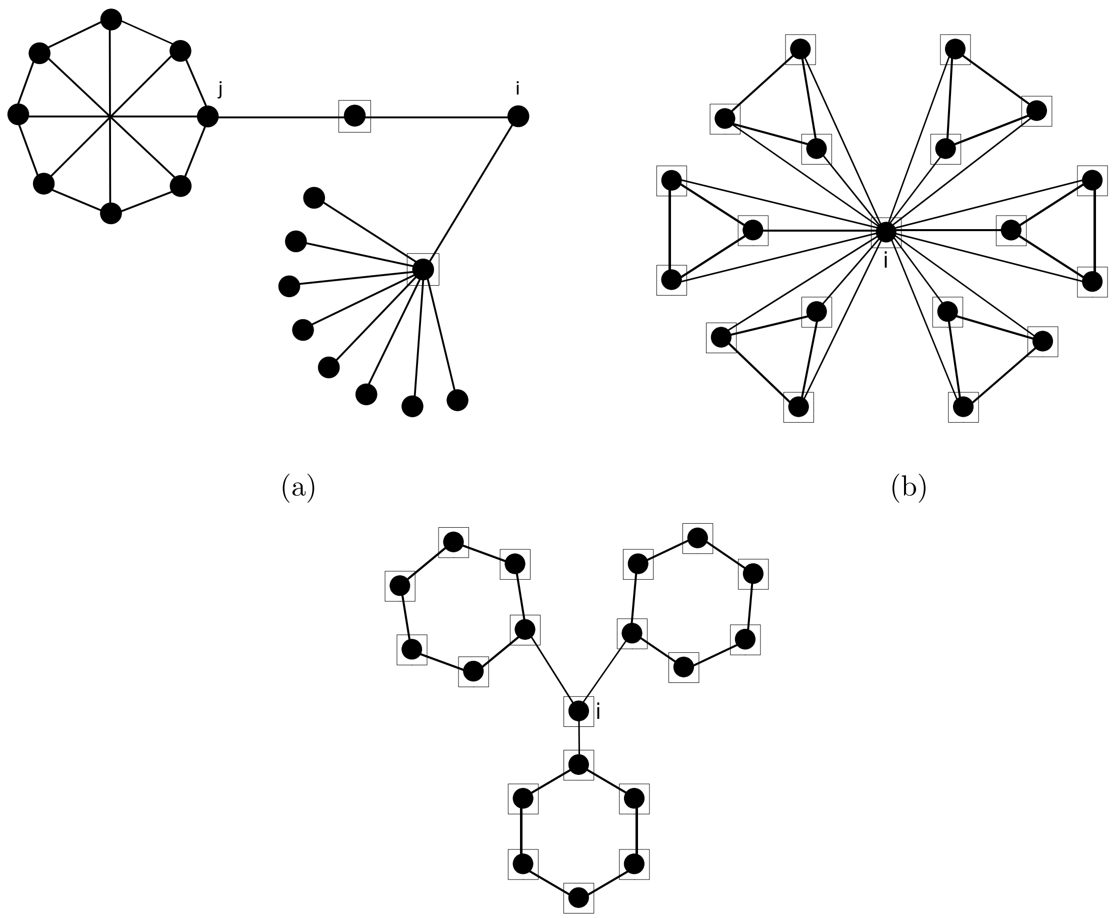

Figure 5: Networks over $n = 19$ nodes with a k-critical node, $k \leq 9$.

6.1 First best outcome

We start the analysis by characterizing the first best. The following definitions will be used. For $B(n) = \{\boldsymbol{b} \in \mathbb{N}^n : b_1 \geq \ldots \geq b_n \geq 0 \text{ and } \sum_i b_i = n\}$, let

$$B^*(n) = \arg\max_{\boldsymbol{b} \in B(n)} \sum_{i=1}^{n} f(b_i)(n - b_i).$$

For $\boldsymbol{b} \in B^*(n)$, we will let $K(\boldsymbol{b})$ denote the maximum i such that b_i is strictly positive. Moreover, let

$$\hat{c}_1(n) = \frac{f(n) - f(n-1)}{n}, \tag{15}$$

$$\hat{c}_2(n) = \frac{f(n) - \frac{1}{n}\sum_{i=1}^{n} f(b_i)(n - b_i)}{n}, \tag{16}$$

$$\hat{c}_3(n) = \frac{f(n) + (n-1)f(n-1)}{n} - \sum_{i=1}^{n} f(b_i)(n - b_i), \tag{17}$$

where $\boldsymbol{b} \in B^*(n)$.

Proposition 7. *Suppose the attack is random and the designer chooses protection as well as design. Then*

(1) If $c \leq \min\{\hat{c}_1(n), \hat{c}_2(n)\}$, the network is connected and all nodes are protected.

(2) If $\hat{c}_1(n) < c \leq \hat{c}_3(n)$, the network is a star and only the centre is protected.

(3) If $c > \max\{\hat{c}_2(n), \hat{c}_3(n)\}$, the network is unprotected and has $K(\boldsymbol{b})$ components, of sizes $b_1, \ldots, b_{K(\boldsymbol{b})}$.

When the first best involves protection, the topologies that are optimal are the same as under intelligent attack. The novel aspect is the structure of the optimal unprotected network. Facing an intelligent threat, there is no point in choosing an unprotected network with a unique largest component; the adversary would remove such a component. Under random attack, the designer may choose an unprotected network with a very large component if the network value function is sufficiently convex. For example, if $f(y) = \alpha^y$, $\alpha \geq e$,[20] then the optimal unprotected network consists of a component of size $(n-1)$ and an isolated node.

6.2 Metcalfe's Law

In this section we present the characterization of optimal networks under random attack when $f(y) = y^2$. As shown in the Appendix, for this value function the optimal unprotected network consists of two components, of sizes $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$. The differences

[20] Where e is the base of the natural logarithm.

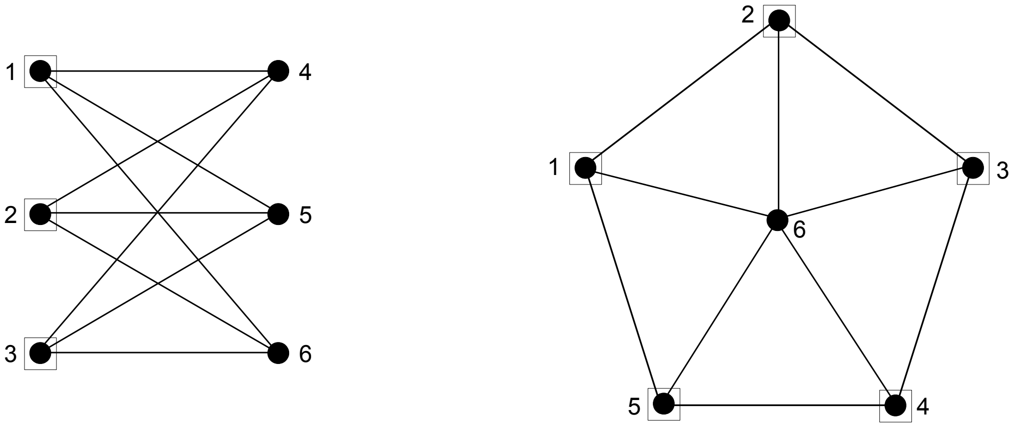

(a) (b)

Figure 6: Networks over $n = 6$ nodes. In the network of Figure 6a, the unprotected nodes $\{4, 5, 6\}$ expose the other nodes to the possibility of contagion, but an attack on any of them neither spreads nor disconnects the network if all other nodes protect. Figure 6b presents an example with only one unprotected node.

in the first best between intelligent and random attack are thus minor in the case of Metcalfe's Law.[21]

Before we state the next result, we need to define the following sets of networks over n nodes:

$$\begin{aligned}\mathcal{G}^{n-u}(N) \quad = \quad & \{G \in \mathcal{G}(N) : \text{ exists } U \subseteq N \text{ such that } |U| = u, \\ & G - \{i\} \text{ is connected for all } i \in U, \text{ and } ij \in G \text{ for } i \in U \text{ iff } j \notin U\}\end{aligned}$$

for $u = 1, \ldots, n-1$. For $u = 0$, let $\mathcal{G}^n(N)$ denote the set of connected networks. To illustrate, suppose that there are $n = 6$ nodes. Figure 6a shows a network in $\mathcal{G}^3(N)$. Note that, e.g., the set of nodes $U = \{4, 5, 6\}$ satisfy the conditions required: they are not linked among themselves, but have links to all other nodes. Moreover, their removal does not disconnect the network. The network in Figure 6a is in $\mathcal{G}^5(N)$. In this case, $U = \{6\}$.

Proposition 8. *Assume $f(y) = y^2$, and suppose the attack is random. If (G, Δ) is a welfare maximising equilibrium, then*

(1) If $c \leq 1$, G is connected and all nodes protect.

(2) If $1 < c \leq \hat{c}_1(n)$, G is in $\mathcal{G}^{n-1}(N)$ and all but one nodes protect.

[21]The minor differences stem, first, from the fact that if n is odd then the optimal unprotected network features a component of size $(n+1)/2$ and another of size $(n-1)/2$. Secondly, under random attack the hub of the center-protected star is attacked with positive probability. This makes the star more attractive, and therefore the threshold for the star to be better than a fully protected network is $\hat{c}_1(n) = \frac{2n-1}{n}$ under random attack, which is smaller than the threshold $c_1(n) = \frac{2n-1}{n-1}$ under intelligent attack. Naturally, this extra benefit of the star under random attack vanishes as n grows.

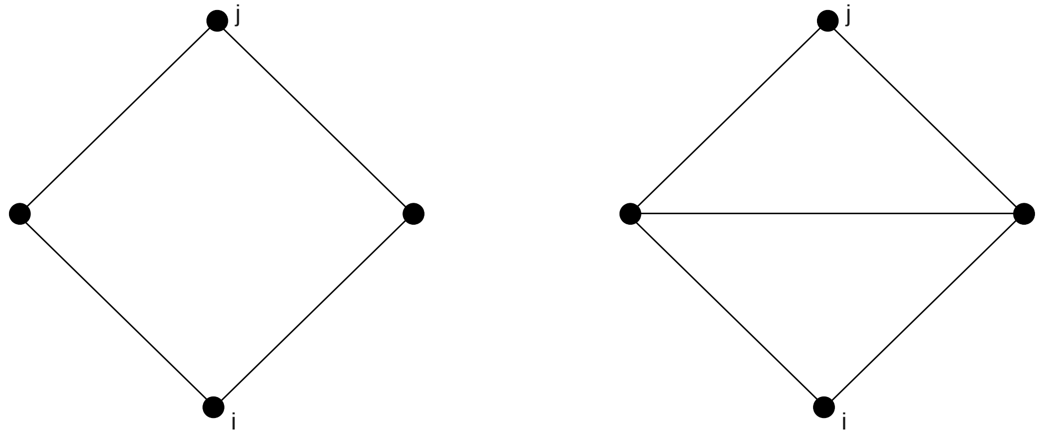

(a) (b)

Figure 7: Let $f(x) = x^2$. Network 7a features strategic complements: node j protecting increases incentives for node i to protect. Network 7b features strategic substitutes: node j protecting decreases incentives for node i to protect.

(3) If $\hat{c}_1(n) < c \leq (n-1) + 1/n$, G is a star and only the centre protects.

(4) If $c > (n-1) + 1/n$, G is unprotected and has two components, of sizes $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$.

Under random failure, investments in security always exhibit positive externalities. Additionally, protection decisions may be strategic substitutes, as well as strategic complements (c.f. Figure 7). The latter possibility is due to the fact that nodes care not only for staying uninfected but also for the benefits they derive from being in the network. Either way, the positive externalities effect always prevails and the over-protection problem is no longer present if the external threat is unintelligent. In effect, the designer will disconnect the network in decentralized equilibrium only for reasons related to under-protection.

Interestingly, when the designer decides to keep the network connected, s/he will not choose any such network (even if nodes coordinate on welfare maximising equilibria!). This also stands in sharp contrast with the case of intelligent attack, where under welfare maximising equilibria the designer could choose any connected network to enforce full protection. For relatively small protection costs, the intelligence of the adversary works for the designer's advantage. Under random attack, the designer needs to choose the network more carefully. In particular, networks that satisfy the properties to belong in $\mathcal{G}^{n-1}(N)$ have equilibria where a subset of nodes are sufficiently exposed so as to secure maximum protection spending in equilibrium.

To illustrate point (2) in Proposition 8, suppose there are $n = 19$ nodes, and that $c = 1.5 \in (1, \hat{c}_1(n))$. Since $c \leq \hat{c}_1(n)$, first best is full protection in a connected network. Note, however, that in any equilibrium of a connected network there will be at least one unprotected node. If all other nodes protect, the individual gain from protection

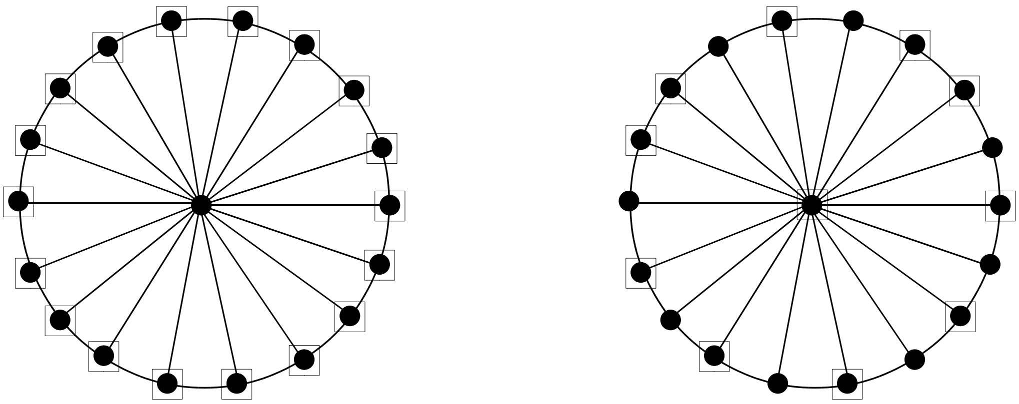

(a) (b)

Figure 8: Two equilibria on the wheel network over $n = 19$ nodes.

is $\frac{1}{n}n - c < 0$. Maximum equilibrium welfare will therefore be achieved if there is a single unprotected node that elimination neither spreads nor disconnects the network. Figure 8a presents the wheel network as an example. Note that protected nodes do not wish to unprotect, for the gain from protection is $\frac{2}{n}n - c > 0$.

Given that multiple protection profiles can be equilibria in a given network, the optimality of the networks presented in Proposition 8 may rely on nodes coordinating on the right equilibrium. If the cost of protection is $c \leq 1$, a node will choose to protect on any network (since $\frac{1}{n}n - c \geq 0$). Moreover, if $\hat{c}_1(n) < c \leq (n-1) + 1/n$, then centre protection is the unique equilibrium of the star network.

What if the cost of protection is small ($c \leq \hat{c}_1(n)$), but not too small ($c > 1$)? We show that in this case a network G attains maximum equilibrium welfare in every equilibrium if and only if G is the complete network. To illustrate this, Figure 8b shows another equilibrium on the wheel network where more than one node is unprotected. The next result characterizes the optimal network under welfare minimizing equilibria.

Proposition 9. *Assume $f(y) = y^2$, and suppose the attack is random. If (G, Δ) is a welfare minimizing equilibrium, then*

(1) If $c \leq 1$, G is connected and all nodes protect.

(2) If $1 < c \leq \hat{c}_1(n)$, G is the complete network and all but one nodes protect.

(3) If $\hat{c}_1(n) < c \leq (n-1) + 1/n$, G is a star and only the centre protects.

(4) If $c > (n-1) + 1/n$, G is unprotected and has two components, of sizes $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$.

The key point is to note that potential coordination problems among nodes are addressed in fundamentally different ways depending on the nature of the attack. Under both intelligent and random attack, the set of networks chosen bearing in mind coordination failures is a strict subset of the possible designs when coordination problems are absent.[22] The reason is that, in both cases, the designer can prevent coordination failures by appropriately choosing the network. However, when the adversary is intelligent, full protection is secured by choosing networks that are *sparse*: there must exist a sufficiently important node that can block the adversary's attack and thus be willing to protect. On the contrary, under random attack maximal protection is achieved by choosing networks that are *dense*: a node must be exposed to an unprotected node, or otherwise it would not have enough incentives to protect; maximal protection in every equilibrium is achieved through maximal exposure, i.e. by designing a complete network.

6.3 General network value function

In this section we discuss in what ways the intuitions brought forward by the case of Metcalfe's Law generalize to other network value functions. Before we state the proposition characterizing welfare maximising equilibria, we need to introduce the following quantities. Let

$$\begin{aligned} t_0(n) &= 0, \\ t_u(n) &= \frac{f(n)}{n^2} + (u-1)\frac{f(n-1)}{n(n-1)}, \text{ for } u = 1, \dots, n. \end{aligned}$$

Proposition 10. *Suppose the attack is random, and let (G, Δ) be a welfare maximising equilibrium of Γ. Then*

(1) G is in $\mathcal{G}^{n-u}(N)$ and exactly u nodes do not protect, if $t_u(n) < c \leq \min\{t_{u+1}(n), \hat{c}_1(n), \hat{c}_2(n)\}$, for $u = 0, \dots, n-1$.

(2) G is a star and only the centre protects, if $\hat{c}_1(n) < c \leq \min\{\hat{c}_3(n), t_n(n)\}$.

(3) G is an optimal unprotected network, if $c > \min\{t_n(n), \max\{\hat{c}_2(n), \hat{c}_3(n)\}\}$.

Recall that under Metcalfe's Law, if full protection is first best ($c \leq \{\hat{c}_1(n), \hat{c}_2(n)\}$) then either all (if $c \leq t_1(n)$) or all except one (if $t_1(n) < c \leq t_2(n)$) nodes protect in equilibrium. The generalization of Proposition 10 shows that this depends on the specific network value function.

Suppose, for example, that there are $n = 6$ nodes and consider again Reed's Law, i.e. $f(y) = 2^y - 1$. The optimal unprotected network consists of two components, of sizes 4 and 2. From this observation, it is straightforward to see that full protection is first best if $c \leq 5.3$ (as $\min\{\hat{c}_1(n), \hat{c}_2(n)\} = \hat{c}_1(n) = 5.3$). Simple calculations indicate that in this

[22] When the aversary is intelligent, we know that the network must be connected and contain a k-critical node with $k < n - c$, and this class of networks is a strict subset of the set of connected networks. Under random attack, the complete network is a strict subset of $\mathcal{G}^{n-1}(N)$.

case $t_4(n) < 5.3 < t_5(n)$. Therefore, when first best is full protection, up to 4 nodes may be unprotected in equilibrium if the cost of protection is large enough. If $t_4(n) < c < 5.3$, in decentralized equilibrium of *any* network there will be at least 4 unprotected nodes. Equilibrium welfare is therefore bounded above by the case where there are exactly 4 unprotected nodes such that an attack on any of them neither spreads nor disconnects the network, and this is only attained by networks in $\mathcal{G}^2(N)$.

As we discussed for the case of Metcalfe's Law, the optimality of the networks that attain maximum equilibrium protection may depend on nodes coordinating on the best equilibrium. For the network value function $f(y) = y^2$ the welfare costs of coordination problems could be avoided by choosing the right topology. We conclude this section with a discussion on whether this finding generalizes to other network value function.

For sufficiently small protection costs, first best networks attain first best payoffs in every decentralized equilibrium for any network value functions.

Fact 1. *Suppose the attack is random and first best involves full protection. If $0 < c \leq t_1(n)$ and G is connected, the unique equilibrium of $\Gamma(G)$ attains first best payoffs.*

If costs of protection are low (so that full defence is first best) but not too low, then every equilibrium on any network features some unprotected nodes. The designer's optimal choice is a network with an equilibrium where the number of unprotected nodes is as small as possible. Suppose $t_1(n) < c \leq t_2(n)$. In the welfare maximising case, the designer chooses a network G with a node l who has a link to all other nodes, and $G - \{l\}$ is connected. There is an equilibrium on such a network where l is the only unprotected node. Notice that the complete network satisfies these properties – l can be any node $i \in N$. The next result states that for any network different from the complete network, a worse equilibrium exists (i.e. one where more than one node does not protect). Thus, the only hope if the designer wants to achieve maximum equilibrium payoffs in every equilibrium is the complete network. However, whether the complete network has only one unprotected node in every equilibrium depends on the network value function.

Fact 2. *Suppose the attack is random, first best involves full protection, and $t_1(n) < c \leq t_2(n)$.*

(1) If G is not the complete network, there exists an equilibrium of $\Gamma(G)$ which does not attain maximum equilibrium welfare.

(2) The complete network G^c attains maximum equilibrium welfare in every equilibrium of $\Gamma(G^c)$ for any $c \leq t_2(n)$ if $\frac{f(n-1)}{(n-1)} \leq u\frac{f(n-u)}{n-u}$ for any $u = 1, \ldots, n-1$.

The intuition for this observation is as follows. Since $t_1(n) < c \leq t_2(n)$, by Proposition 10 we know that there are at least one unprotected node in every equilibrium. This is true because for $c > t_1(n)$ any node who has all its neighbours protected prefers not to protect. For any network that is not the complete network, we can construct an worse equilibrium, where there are two unprotected nodes. This is what is established in point (1) of Fact 2.

By creating maximal exposure, there are no equilibria in the complete network where exactly two nodes do not protect. But nodes may 'get stuck' in worse equilibria in the complete network, and whether this is possible depends on the network value function. Increasing the probability of contagion by creating exposure taps the substitutes aspect of protection decisions. But nodes value being connected to surviving individuals – this is the complements aspect of protection decisions. Thus, if few nodes are protected it may not pay off to protect. Consider, for example, $f(y) = 2^y - 1$ with $n = 6$ nodes, and suppose that $c = 2.6 \in (t_1(n), t_2(n))$. By Proposition 10, the complete network has an equilibrium with full protection. It also has, however, an equilibrium with no protection.[23] The condition in point (2) of Fact 2 bounds the convexity of the value function, and thus bounds the strength of complementarities in protection. This condition holds, for example, if $f(y) = y^a$, $a \leq 2$.

7 Concluding remarks

In this paper we studied the problem of mitigating inefficiencies resulting from protection decentralization by appropriate network design.

Motivated by the example of cybersecurity, we first took up the case of an intelligent threat. An efficient equilibrium may exhibit too much or too little investment in security. The problem of over-protection problem arises for intermediate costs of protection, and is best addressed by disconnecting the network into unequal components, and sacrificing some nodes. The problem of under-protection is more standard and reflects the public good aspect of security. It arises at larger costs of security and is addressed by creating networks with equal components. Finally, inefficient equilibria arise due to strategic complementarity in security. They are addressed by creating networks that are 'sparse' and contain 'critical' nodes. This sparseness gives rise to nodes that can prevent attacks from spreading, and thus save large parts of the network. Although the first best cannot be attained when over-protection pressures prevail, network design puts a bound on the welfare costs of decentralization.

Finally, motivated by problems in epidemiology, we studied optimal design in the face of random attack. The over-protection problem is no longer present, whereas under-protection problems may be mitigated in a diametrically opposite way: namely, by creating dense networks that expose the individuals to the risk of contagion.

[23] If no other node protects in the complete network, then a node's gain from protecting equals $\frac{1}{n}\frac{f(n)}{n} + \frac{n-1}{n}f(1) - c = 2.58 - c < 0$. Therefore, no protection is an equilibrium. In fact, if no protection is an equilibrium defence on the complete network, then it is an equilibrium defence on any d-connected network, $d \geq 2$. The optimal network in these case must therefore be 1-connected.

References

D. Acemoglu, A. Malekian, and A. Ozdaglar. Network security and contagion. *MIT Mimeo*, 2013.

D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.

J. Aspnes, K. Chang, and A. Yampolskiy. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *Journal of Computer and System Sciences*, 72(6):1077–1093, 2006.

M. Baccara and H. Bar-Isaac. How to organize crime? *Review of Economic Studies*, 75 (4):1039–1067, 2008.

L. Blume, D. Easley, J. Kleinberg, R. Kleinberg, and E. Tardos. Network formation in the presence of contagious risk. *ACM Transactions on Economics and Computation*, 1 (2):6:1–6:20, 2013.

B. Bollobás. *Modern Graph Theory*. Springer, July 1998.

A. Cabrales, P. Gottardi, and F. Vega-Redondo. Risk sharing and contagion in networks, 2013. Working Paper.

M. Elliot, B. Golub, and M. O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):315–353, 2014.

S. Goyal and A. Vigier. Attack, Defence and Contagion in Networks. *Review of Economic Studies*, Forthcoming, 2014.

H. Kunreuther and G. Heal. Interdependent Security. *The Journal of Risk and Uncertainty*, 26:231–249, 2003.

T. Moore, R. Clayton, and A. R. The economics of online crime. *Journal of Economic Perspectives*, 23(3):3–20, 2009.

A Equilibrium existence

Proof of Lemma 1. Let x be a strategy of A in $\Gamma(G)$ such that for all $\Delta \subseteq N$, $x(\Delta)$ is a best response to Δ. Network G and strategy x define game $\Gamma(G, x)$ with set of players N such that, given defence Δ induced by a strategy profile of the nodes $(\delta_1, \ldots, \delta_n)$, the utility of player i is $\tilde{u}^i(\Delta) = U^i(G, \Delta, x(\Delta))$.
We will show that $\Gamma(G, x)$ has a Nash equilibrium. To show that we will construct a set of defended nodes, Δ^*, such that the corresponding strategy profile of the nodes is Nash equilibrium of $\Gamma(G, x)$.
There are two cases possible. First, suppose that for all components $C \in \mathcal{C}(G)$, $\frac{f(|C|)}{|C|} \geq c$. I this case $\Delta^* = N$ is an equilibrium of $\Gamma(G, x)$, as any node that would deviate and drop protection, would obtain payoff $0 \leq \frac{f(|C|)}{|C|} - c$.
Second, suppose that there exists $C \in \mathcal{C}(G)$ such that $\frac{f(|C|)}{|C|} < c$. Let $\mathcal{A}(G|c) = \{C \in \mathcal{C}(G) : f(|C|)/|C| < c\}$ be the set of all such components. We construct Δ^* using the following algorithm.

- $\Delta^* := N \setminus \bigcup \mathcal{A}(G|c)$, i.e. Δ^* protects all the nodes in components where protection yields non-negative payoffs to the protected nodes; for any $C \in \mathcal{A}(G|c)$, $C \cap \Delta^* = \varnothing$; note that $x(\Delta^*)$ removes $C \in \mathcal{A}(G|c)$ of maximal size.
- While there exists $i \in \Delta^*$ such that $x(\Delta^* \setminus \{i\}) \in \mathcal{A}(G|c)$ do
  - $\Delta^* := \Delta^* \setminus \{i\}$.

Clearly the algorithm stops, as in every step at least one node is removed from Δ^*. Moreover, $x(\Delta^*)$ removes $C \in \mathcal{A}(G|c)$ of maximal size and no node in C has incentive to protect. The algorithm ensures that no node in Δ^* has incentive to drop protection either. Hence Δ^* is an equilibrium protection of $\Gamma(G, x)$ and (Δ^*, x) is an equilibrium of $\Gamma(G)$. ◻

B The First Best Outcome

Proof of Proposition 1. Let (G, Δ) be a first best protected network. Three cases are possible.

Case (i). $\Delta = N$ Clearly in this case G must be a connected network.

Case (ii). $\varnothing \subsetneq \Delta \subsetneq N$ In this case **A** removes at least one node from G and so gross welfare is bounded from above by $f(n-1)$. Star network is the unique network that attains this upper bound by using only one unit of protection. This is the lowest possible number of protected nodes possible in Case (ii). Thus G is a star and $\Delta = \{i\}$, where i is is the centre of G.

Case (iii). $\Delta = \varnothing$ As long as $n > 1$, any disconnected G yields higher welfare than a connected network in this case. Moreover, there are at most two sizes of components in $\mathcal{C}(G)$. For assume otherwise and let $C_1, C_2, C_3 \in \mathcal{C}(G)$ be such that $|C_1| > |C_2| > |C_3|$. Then, since f is strictly increasing and strictly convex, **D** is better of by moving a node from C_3 to C_2. Lastly, if C_1 is the component of maximal size in $\mathcal{C}(G)$, then there is at most one component $C \in \mathcal{C}(G)$ with $|C| < |C_1|$. If there was another component $C' \in \mathcal{C}(G)$ with $|C'| = |C|$, then, since f is strictly increasing and strictly convex, **D** would be better off by moving a anode from C' to C. It is straightforward to see that the number and the sizes of the components are as stated in the proposition.
Comparing the payoffs of (i)-(iii) yields the desired result. □

C The price of decentralization

To prove Lemma 2, we will use the following two results. Let $\mathcal{G}^{full}(N,c)$ denote the set of connected networks such that all nodes protect in any equilibrium. That is,

$$\mathcal{G}^{full}(N,c) = \{G \in \mathcal{G}(N) : G \text{ is connected and } \Delta(G) = N \text{ for any } (\Delta, x) \in \mathcal{E}(c|G)\}.$$

Lemma 5. *$\mathcal{G}^{full}(N,c) \neq \varnothing$ if and only if $c \leq c_0(n)$.*

Proof. For left to right implication, note that if $c > c_0(n) = \frac{f(n-1)}{n-1}$, then any node strictly prefers to protect only if all other nodes survive. Therefore, for any network G, the strategy profile where no node protects is an equilibrium of $\Gamma(G)$.
For right to left implication, suppose that $c \leq c_0(n) = \frac{f(n-1)}{n-1}$ and let G be a star network and i be the centre of G. Take any equilibrium (Δ, x) of $\Gamma(G)$. It must be that $i \in \Delta$ as otherwise i would be removed by **A** obtaining payoff 0 instead of $\frac{f(n-1)}{n-1} - c \geq 0$. Similarly, if there is $j \in N \setminus \{i\}$ such that $j \notin \Delta$, then $k = x(\Delta) \notin \Delta$ and k is better off by protecting, which yields payoff at least $\frac{f(n-1)}{n-1} - c \geq 0$. □

Fact 3. *For all $q \in Q^*(n)$, $\frac{f\left(\left\lfloor \frac{n}{q-1} \right\rfloor\right)}{\left\lfloor \frac{n}{q-1} \right\rfloor} < c_3(n)$.*

Proof of Lemma 2. Consider first $c \leq \min\{c_0(n), c_1(n), c_2(n)\}$. Since $c \leq \min\{c_1(n), c_2(n)\}$, first best is attained through full protection in a connected network. Since $c \leq c_0(n)$, by Lemma 5 there exists connected G such that $\Delta = N$ in every equilibrium of $\Gamma(G)$.
Consider next $c > \max\{c_2(n), c_3(n)\}$. By Fact 3, this implies that $c > \frac{f\left(\left\lfloor \frac{n}{q-1} \right\rfloor\right)}{\left\lfloor \frac{n}{q-1} \right\rfloor}$. Therefore, if $c > \max\{c_2(n), c_3(n)\}$ and G is a first best network, $\Delta = \varnothing$ in every equilibrium of $\Gamma(G)$. □

Proof of Lemma 3. Let G be a star network. For a contradiction, suppose there is an equilibrium (Δ, x) on G with $\Delta \subsetneq N$.
Suppose that $\frac{f(n-1)}{n-1} < c$. In this case any protected node $i \in \Delta$ would be better off by deviating to no protection, as **A** removes at least one node from G and so the payoff to i, $U^i(G, \Delta, x(\Delta)) \leq \frac{f(n-1)}{n-1} - c < 0$.
Suppose that $\frac{f(n-1)}{n-1} \geq c$ and let i be the centre of star G. If i protects, it obtains at least the payoff of $\frac{f(n-1)}{n-1} - c \geq 0$. Hence i prefers to protect, regardless of protection decisions of other nodes. Let $j = x(\Delta)$ be the node attacked by **A**. As in the case of i, j prefers to protect, a contradiction with the assumption that (Δ, x) is an equilibrium. □

Proof of Theorem 1. Point 1 follows directly from the discussion in the main text. For point 2 notice that, by Lemma 5, **D** can enforce full protection by choosing the right connected network if and only if $\frac{f(n-1)}{n-1} \geq c$.
Since $\frac{f(y)}{y}$ is increasing, either $\frac{f(n-1)}{n-1} \geq c$ for sufficiently large n, or $\frac{f(n-1)}{n-1} < c$ for all $n \geq 1$.

Consider the former case. Then, for sufficiently large n, **D** could choose a connected network where all nodes protect in every equilibrium. This attains a welfare of $f(n) - nc$. Thus

$$\lim_{n\to\infty} \mathrm{PoA}(c,n) \le \lim_{n\to\infty} \frac{f(n-1)-c}{f(n)-nc} = \lim_{n\to\infty} \frac{1 - \frac{n}{f(n)}\frac{f(n)-f(n-1)}{n} - \frac{c}{f(n)}}{1-\frac{n}{f(n)}c}. \tag{18}$$

Suppose that $\frac{f(n)}{n}$ is unbounded, i.e. $\lim_{n\to\infty}\frac{f(n)}{n} = +\infty$. Then, by the fact that $\frac{f(y)-f(y-1)}{y-1}$ is bounded,

$$\lim_{n\to\infty} \frac{1 - \frac{n}{f(n)}\frac{f(n)-f(n-1)}{n} - \frac{c}{f(n)}}{1-\frac{n}{f(n)}c} = 1. \tag{19}$$

Suppose now that $\lim_{n\to\infty}\frac{f(n)}{n} = p < +\infty$. In this case $\lim_{n\to\infty}\frac{f(n)-f(n-1)}{n-1} = 0$ and

$$\lim_{n\to\infty} \frac{1 - \frac{n}{f(n)}\frac{f(n)-f(n-1)}{n} - \frac{c}{f(n)}}{1-\frac{n}{f(n)}c} = \frac{p}{p-c}. \tag{20}$$

Assume now that $\frac{f(n-1)}{n-1} < c$, for all $n \ge 2$. If **D** chooses the fully disconnected network, then

$$\lim_{n\to\infty} \mathrm{PoA}(c,n) \le \lim_{n\to\infty} \frac{f(n-1)-c}{(n-1)f(1)} = \frac{p}{f(1)}. \tag{21}$$

This completes the proof. □

D Decentralized security and optimal design

We structure the proofs of optimal design in the following way. Results that are for general network value function are shown first. Based on these results, we then show the results for specific functional forms. We start with welfare maximising equilibria, and then move to welfare minimizing equilibria.

D.1 Welfare maximising equilibria

To prove Proposition 4, we need a couple of intermediate results and some notation. Given a network $G \in \mathcal{G}(N)$ and a set of nodes $U \subseteq N$, the *neighbourhood* of U in G is the set of nodes $\partial_G(U) = \{j \in N \setminus U : \exists i \in U.\ ij \in E\}$. In the case of a singleton set $\{i\}$, its neighbourhood in G is the set of *neighbours* of i in G. In this case we omit the curly brackets and write $\partial_G(i)$ rather than $\partial_G(\{i\})$. Given network G and set of protected nodes Δ, we will say that a component $C \in \mathcal{C}(G)$ is *partially protected* under Δ if $C \cap \Delta \neq \varnothing$ and $C \setminus \Delta = \varnothing$.

Lemma 6. *Let $G \in \mathcal{G}(N)$. In every equilibrium (Δ, x) of $\Gamma(G)$, if two protected nodes are connected in network G, then they are connected in the residual network $G - E_{x(G,\Delta(G))}(G|\Delta)$.*

Proof. Assume otherwise and let $i, j \in \Delta$ be connected in G and disconnected in $G - E_{x(\Delta)}(G|\Delta)$. Then, under defence Δ all paths between i and j go through $E_{x(\Delta)}(G|\Delta)$. Pick $l \in E_{x(\Delta)}(G|\Delta)$ such that l is on a path from i to j in G and is a neighbour of $C_i(G[\Delta(G)])$. Let $\Delta' = \Delta \cup \{l\}$.
It must be that $x(\Delta') \in C_i(G)$ or l is strictly better off by getting positive payoff (it gets higher payoff than its protected neighbour gets under Δ). Thus, there must be at least another unprotected node, $l' \in C_i(G)$. There are two cases possible:

Case (i). $x(\Delta')$ is reachable from l in $G - \Delta$. In this case $E_{x(\Delta')}(G|\Delta') \subseteq E_{x(\Delta)}(G|\Delta)$ and l gets strictly higher payoffs than its protected neighbour did under original strategy profile; thus l is strictly better off with payoff > 0 (a contradiction).

Case (ii). $x(\Delta')$ is not reachable from l in $G - \Delta$. In this case, i and j are connected in $G - E_{x(\Delta')}(G|\Delta')$. Suppose that $x(\Delta') \notin C_i(G - E_{x(\Delta)}(G|\Delta))$ (the case with C_j is analogous). Then $|C_i(G - E_{x(\Delta')}(G|\Delta'))| > |C_i(G - E_{x(\Delta)}(G|\Delta))|$ (the component gets extended by node j, at least). Since $l \in C_i(G - E_{x(\Delta')}(G|\Delta'))$ so l is strictly better off (a contradiction). □

The next lemma characterizes the structure that a partially protected component must have for it to be attacked in equilibrium.

Lemma 7. *Suppose that (G, Δ, x) is a welfare maximising equilibrium of Γ. Let $\Delta^* = \Delta(G)$, $x^* = x(G, \Delta^*)$ and $X^* = E_{x^*}(G|\Delta^*)$. If x attacks a partially protected component $P \in \mathcal{C}(G^*)$, then there exists a unique set of nodes $Y \neq \varnothing$ such that*

1. $G[Y]$ is connected.

2. $\Delta^* \cap P = \partial_G(Y)$, i.e. the neighbourhood of Y is the set of protected nodes in P.

3. $|\mathcal{C}(G[P] - Y)| = |\Delta^* \cap P| \geq 2$ and for each $C \in \mathcal{C}(G[P] - Y)$, $|C \cap \Delta^*| = 1$.

Moreover, for all $i \in \partial_G(\Delta^*) \cap X^*$, $x(G, \Delta^* \cup \{i\}) \in Y$.

Proof. Let P be a partially protected components attacked by x, as stated in the lemma. Let $\Delta^i = \Delta^* \cup \{i\}$, $x^i = x(G, \Delta^i)$ be equilibrium response of **A** to G and defence Δ extended with a node $i \in N$, and let $X^i = E_{x^i}(G|\Delta^i)$ be the set of nodes eliminated by x^i. We prove the lemma in the following three steps.

Step 1. For all $i \in \partial_G(\Delta) \cap X$, $|\mathcal{C}(G[P] - X^i)| \geq 2$. By Lemma 6, the residual component $G[P] - X$ is connected and its value is $f(|P| - |X|)$. Let $i \in \partial_G(\Delta) \cap X$ be a node removed by attack x and neighbouring a protected node. Suppose, to the contrary, that $|\mathcal{C}(G - X^i)| \leq |\mathcal{C}(G)|$, that is component $G[P]$ does not get disconnected by X^i. If so, then $|X^i| \leq |X|$ (as otherwise x^i would be a better response to (G, Δ^*) than x) and so the value of residual component $G[P] - X^i$, $f(|P| - |X^i|) \geq f(|P| - |X|)$. Payoff to i in $G[P] - X^i$, when it protects, is $f(|P| - |X^i|)/(|P| - |X^i|) - c \geq f(|P| - |X|)/(|P| - |X|) - c$ (as f is increasing and convex). Since $f(|P| - |X|)/(|P| - |X|) - c \geq 0$ (as there are protected nodes in $P \cap \Delta^*$ that get exactly this payoff) so i is better off. Thus $G[P]$ must get disconnected by X^i, i.e. $|\mathcal{C}(G[P] - X^i)| \geq 2$.

Step 2. For all $i \in \partial_G(\Delta^*) \cap X^*$ and any $C \in \mathcal{C}(G[P] - X^i)$, $|C \cap \Delta^*| = 1$. Pick any $i \in \partial_G(\Delta^*) \cap X^*$ and any $C \in \mathcal{C}(G[P] - X^i)$. Clearly it must be that $|C \cap \Delta| \geq 1$, as otherwise C would be removed by x^i. Let $j \in \partial_G(X^i) \cap C$ (any node in $\partial_G(X^i)$ must be protected). Let $G' = (N, E')$ be a network obtained from G by removing all links to nodes from $C \setminus \{j\}$ and linking these nodes to j only, i.e. $E' = (E \setminus E[C \setminus \{j\}]) \cup \{jl : l \in C \setminus \{j\}\}$; additionally, in the case of $j \in \partial_G(i)$, all nodes from X^* are linked to form a clique, i.e. $E' = (E \setminus E[C \setminus \{j\}]) \cup \{jl : l \in C \setminus \{j\}\} \cup \{lr : l, r \in X, l \neq r\}$.

Consider a strategy profile (Δ', x') in $\Gamma(G')$ such that

- $\Delta' = (\Delta^* \setminus C) \cup \{j\}$.
- $x'(\Delta') = x^*$.
- For all $l \in \Delta'$, $x'(\Delta' \setminus \{l\}) = x(G, \Delta' \setminus \{l\})$.
- For all $l \in P \setminus (\Delta' \cup C)$, $x'(\Delta' \cup \{l\}) = x(G, \Delta \cup \{l\})$ (note that $P \setminus (\Delta' \cup C) = P \setminus (\Delta \cup C)$).
- For all $l \in (C \setminus \{j\})$, $x'(\Delta' \cup \{l\}) = x^*$.

The strategy profile (Δ', x') is an equilibrium of $\Gamma(G')$. The responses of A to Δ' and any single node deviations from Δ' are best responses, because they are best responses to (G, Δ) and any single node deviations from $\Delta' \subseteq \Delta^*$. None of the nodes is better off by deviating from its strategy, as they obtain the same payoffs as under $(\Delta^*, x(G, \cdot))$ in the game $\Gamma(G)$.

Since (Δ', x') is an equilibrium of $\Gamma(g')$ and it yields a better welfare than $(\Delta^*, x(G, \cdot))$ in $\Gamma(G)$ (as less defence resources are used), so we get a contradiction with the assumption that (G, Δ, x) is an equilibrium. Hence it must be that $|C \cap \Delta| = 1$.

Step 3. The structure of $G[P]$. Clearly $G[X^i]$ is connected. Moreover, since for any component $C \in \mathcal{C}(G[P] - X^i)$, $|C \cap \Delta^*| = 1$, so $\Delta^* \cap P = \partial_G(X^i)$, i.e. the neighbourhood of X^i is the set of protected nodes. Additionally, by Step 1, $\mathcal{C}(G[P] - X^i) \geq 2$. Thus taking $Y = X^i$ we have a set of nodes that satisfies points $1 - 3$ stated in the lemma. It remains to be shown that it is unique such set of nodes. Assume to the contrary that there is a set of node $Y' \neq Y$ that satisfies points $1 - 3$ as well. It cannot be that $Y' \cap Y \neq \varnothing$, because $\partial_G(Y')$ would contain unprotected nodes from Y (which violates point 2 for Y'). But then, by point 2 for Y', $G[\Delta^* \cap P] - Y$ is connected, which violates point 3 for Y.

Uniqueness of Y together with points $1 - 3$ and Step 1 imply that for all $i \in \partial_G(\Delta^*) \cap X$, $x(G, \Delta \cup \{i\}) \in Y$. □

Let (G, Δ, x), $\Delta^* = \Delta(G)$, $x^* = x(G, \Delta^*)$, $P \in \mathcal{C}(G)$, and $X^* = E_{x^*}(G|\Delta^*)$ be as defined in Lemma 7. Let Y be the set of nodes satisfying points $1 - 3$ of Lemma 7. Suppose that $\Delta^* \cap P = \{j_0, \ldots, j_d\}$ with $j_0 \in \partial_G(X^*)$. Let $\mathcal{C}(G[P] - Y) = \{Z_0 \cup X, Z_1, \ldots, Z_d\}$ with $j_i \in Z_i$, for all $i \in \{1, \ldots, d\}$ (by Lemma 7 this is possible; in the component of $G[P] - Y$ containing j_0 we distinguish two subsets: X and the remaining set of nodes Z_0). The structure of $G[P]$ is illustrated in Figure 4.

The components of G can be divided into three disjoint sets (some of them possibly empty): $\{P\}$, $\mathcal{D} = \{C \in \mathcal{C}(G) : C \cap x^* = \varnothing \text{ and } C \cap \Delta^* \neq \varnothing\}$ (the set of not attacked components, protected under Δ^*), and $\mathcal{U} = \{C \in \mathcal{C}(G) : C \cap (x^* \cup \Delta^*) = \varnothing\}$ (the set of not attacked components not protected under Δ^*).

In the following lemmas we establish further properties of network G and subnetwork $G[P]$.

Lemma 8.

$$f(|X|) \geq f(|P| - |X|) - f(|P| - |X| - 1). \tag{22}$$

Proof. Assume to the contrary that

$$f(|X|) < f(|P| - |X|) - f(|P| - |X| - 1). \tag{23}$$

Since each node j_i, with $i \in \{1, \dots, d\}$ is protected, so if the node would not protect, the adversary would remove it. Thus for all $i \in \{1, \dots, d\}$,

$$f(|P|) - f(|X| + |Z_0|) - \sum_{i=1, i \neq q}^{d} f(|Z_i|) > f(|P|) - f(|Z_0| + |Y| + \sum_{i=1}^{d} |Z_i|) \tag{24}$$

which implies

$$f(|X| + |Z_0|) < f(|Z_0| + |Y| + \sum_{i=1}^{d} |Z_i|) - \sum_{i=1, i \neq p}^{d} f(|Z_i|) \tag{25}$$

and further

$$f(|X| + |Z_0|) < f(|Z_0| + |Y| + \sum_{i=1}^{d} |Z_i|). \tag{26}$$

By the fact that f is strictly increasing, it follows that

$$|X| < |Y| + \sum_{i=1}^{d} |Z_i|. \tag{27}$$

On the other hand, since removing the nodes from X is better than attacking a node in Y and disconnecting the component, we have

$$f(|P|) - f(|Z_0| + |Y| + \sum_{i=1}^{d} |Z_i|) > f(|P|) - f(|X| + |Z_0|) - \sum_{i=1}^{d} f(|Z_i|), \tag{28}$$

which implies

$$f(|X| + |Z_0|) + \sum_{i=1}^{d} f(|Z_i|) > f(|Z_0| + |Y| + \sum_{i=1}^{d} |Z_i|). \tag{29}$$

Since $|P| = |X| + |Z_0| + |Y| + \sum_{i=1}^{d} |Z_i|$, so Equation (23) implies

$$f(|Z_0| + |Y| + \sum_{i=1}^{d} |Z_i|) > f(|X|) + f(|Z_0| + |Y| - 1 + \sum_{i=1}^{d} |Z_i|). \tag{30}$$

This, together with Equation (29) implies

$$f(|X| + |Z_0|) + \sum_{i=1}^{d} f(|Z_i|) > f(|X|) + f(|Z_0| + |Y| - 1 + \sum_{i=1}^{d} |Z_i|), \tag{31}$$

from which we get

$$f(|X| + |Z_0|) - f(|X|) > f(|Z_0| + |Y| - 1 + \sum_{i=1}^{d} |Z_i|) - \sum_{i=1}^{d} f(|Z_i|) \tag{32}$$

and further, by convexity of f,

$$f(|X| + |Z_0|) - f(|X|) > f(|Z_0| + |Y| - 1 + \sum_{i=1}^{d} |Z_i|) - f(\sum_{i=1}^{d} |Z_i|). \tag{33}$$

and

$$f(|X| + |Z_0|) - f(|X|) > f(|Z_0| + |Y| - 1 + \sum_{i=1}^{d} |Z_i|) - f(|Y| - 1 + \sum_{i=1}^{d} |Z_i|), \quad (34)$$

as f is strictly increasing and $|Y| \geq 1$. Since f is strictly increasing and strictly convex, this yields

$$|X| > |Y| - 1 + \sum_{i=1}^{d} |Z_i| \quad (35)$$

and further, by the fact that $|X|$, $|Y|$ and $|Z_1|, \dots, |Z_d|$ are integers,

$$|X| \geq |Y| + \sum_{i=1}^{d} |Z_i|, \quad (36)$$

a contradiction with Equation (27). Thus we have shown that it must be that $f(|X|) \geq f(|P| - |X|) - f(|P| - |X| - 1)$. □

Fact 4.

$$\frac{f(|X| + |Z_0|)}{|X| + |Z_0|} \leq c \leq \frac{f(|P| - |X|)}{|P| - |X|}. \quad (37)$$

Proof. If it was $f(|X| + |Z_0|)/(|X| + |Z_0|) > c$, then it would be profitable for a node $i \in \partial_{g^*}(j_0) \cap X$ to protect. If it was $c > \frac{f(|P|-|X|)}{|P|-|X|}$, then it would be profitable for any node $i \in P \cap \Delta$ not to protect. □

Corollary 2.

$$2|X| + |Z_0| \leq |P|. \quad (38)$$

Proof. Since $f(y)$ is strictly increasing and strictly convex, so $f(y)/y$ is strictly increasing. Thus, by Equation (37), $|X| + |Z_0| \leq |P| - |X|$ and Equation (38) follows. □

As a corollary from Lemma 8, Fact 4 and Corollary 2 we get that G must have at least one not attacked component, which implies Proposition 4.

Proof of Proposition 4. We prove the proposition by showing that $\mathcal{D} \cup \mathcal{U} \neq \varnothing$.
Assume, to the contrary, that $\mathcal{C}(G) = \{P\}$. Let G' be a network consisting of two components, a clique over the set of nodes X and a star over the set of nodes $V \setminus X$, with centre i. Consider the strategy profile (Δ', x') of the game $\Gamma(G')$ with $\Delta' = \{i\}$, $x'(G', \Delta') \in X$ and $x'(G', \Delta'')$ being a best response to (G', Δ''), for $\Delta'' \neq \Delta'$. Strategy profile (Δ', x') is an equilibrium of game $\Gamma(G')$: by Equation (22), $x'(G', \Delta')$ is a best response to (G', Δ'); by Equation (37), none of the nodes in X can be better off by choosing protection, while being protected in $P \setminus X$ yields non-negative payoff; by Equation (38, **A** would attack $G[P \setminus X]$ if i did not protect (recall that $|Z_0| \geq 1$ as $j_0 \in Z_0$).
Since (Δ', x') is an equilibrium of $\Gamma(G')$ so G' yields a strictly better payoff to **D** than G, a contradiction with the assumption that (G, Δ, x) is a welfare maximising equilibrium. Thus it must be that $\mathcal{D} \cup \mathcal{U} \neq \varnothing$. □

To prove Proposition 2, we need three intermediate steps. Lemma 9 shows that, for any f, if the network is not connected in a welfare maximising equilibrium, then there is no fully protected component. Lemma 10 shows for $f(y) = y^2$ that if G is not connected and the adversary attacks a protected component, then there exists another protected component in G. Based on this intermediate result, Lemma 11 shows that if G is not connected the adversary does not attack a protected component.

Lemma 9. *Let (G, Δ, x) be a welfare maximising equilibrium. If G is not connected, then there is no fully protected component.*

Proof. Suppose G is not connected, and there exists component $X \in \mathcal{C}(G)$ such that $X \subseteq \Delta$. Clearly, X must be the only fully protected component, or otherwise **D** would be strictly better off by merging all fully protected components.

We compare two modifications to G. Network G' is obtained by attaching another component $Y \in \mathcal{C}(G)$ to X, where $x(G, \Delta) \notin Y$ if possible. We present the case where $x(G, \Delta) \notin Y$; if G has only two components the proof is analogous. First, note that there exists an equilibrium of $\Gamma(G')$ were all nodes in $X \cup Y$ protect and **A** attacks the same unprotected node or attacks a protected node in $X \cup Y$. For G to be optimal, it must be that G' does not attain higher welfare. Let $|X| = |X''| + |Y|$, and denote with $p_Y < |Y|$ the number of nodes protected in Y. Then,

$$\begin{aligned} f(|X''| + 2|Y|) - (|X''| + 2|Y|)c &\leq f(|X''| + |Y|) + f(|Y|) - (|X''| + |Y| + p_Y)c \\ \Leftrightarrow c(|Y| - p_Y) &\geq f(|X''| + 2|Y|) - f(|X''| + |Y|) - f(|Y|). \end{aligned} \tag{39}$$

The second modification consists of network G'', formed from network G as follows. Change X into a star, and detach $|Y|$ spokes from it to form a copy of Y. The nodes that have not been detached from X form a component X''. Let Δ'' denote the equilibrium defence profile in $\Gamma(G'')$.

Case (i). $X'' \subseteq \Delta$ For G to be optimal, it must be that

$$\begin{aligned} f(|X''|) + 2f(|Y|) - (2p_Y + |X''|)c &\leq f(|X''| + |Y|) + f(|Y|) - (p_Y + |X''| + |Y|)c \\ \Leftrightarrow c(|Y| - p_Y) &\leq f(|X''| + |Y|) - f(|X''|) - f(|Y|). \end{aligned}$$

Combining this condition with (39), we have that $f(|X''| + 2|Y|) - f(|X''| + |Y|) \leq f(|X''| + |Y|) - f(|X''|)$, which contradicts f being convex.

Case (ii). $X'' \subsetneq \Delta$ and $x(G'', \Delta'') \notin X''$ Following the same steps as in Case (i) leads to a contradiction.

Case (iii). $X'' \cap \Delta = \varnothing$ and $x(G'', \Delta'') \notin X''$ Following the same steps as in Case (i) leads to a contradiction.

Case (iv). $X'' \cap \Delta = \varnothing$ and $x(G'', \Delta'') \in X''$ Since nodes in X'' are eliminated, it must be that $c > \frac{f(|X''|)}{|X''|}$, or

$$|X''|c > f(|X''|). \tag{40}$$

Let Z denote the component attacked in equilibrium (G, Δ, x). By Lemma 6, the payoff to the designer from this component in the original network is $f(|Z| - |E_{x(G|\Delta)}(G|\Delta)|)$. Then, for G to be optimal, it must be that

$$f(|X''| + |Y|) + f(|Z| - |E_{x(G|\Delta)}(G|\Delta)|) - (|X''| + |Y|)c \geq f(|Y|) + f(|Z|) - p_Y c,$$

or, equivalently,

$$|X''|c + (|Y| - p)c \leq f(|X''| + |Y|) - f(|Y|) - [f(|Z|) - f(|Z| - |E_{x(G|\Delta)}(G|\Delta)|)].$$

We can combine this condition with (40) to obtain:

$$f(|X''|) + (|Y| - p_Y)c < f(|X''| + |Y|) - f(|Y|) - [f(|Z|) - f(|Z| - |E_{x(G|\Delta)}(G|\Delta)|)].$$

Rearranging yields:

$$(|Y| - p_Y)c < f(|X''| + |Y|) - f(|X''|) - f(|Y|) - [f(|Z|) - f(|Z| - |E_{x(G|\Delta)}(G|\Delta)|)]. \tag{41}$$

For G to be optimal, c must be such that (39) and (41) holds. Thus, it must be that

$$\begin{aligned} f(|X''| + 2|Y|) - f(|X''| + |Y|) &< f(|X''| + |Y|) - f(|X''|) \\ &\quad -[f(|Z|) - f(|Z| - |E_{x(G|\Delta)}(G|\Delta)|)] \\ &\leq f(|X''| + |Y|) - f(|X''|), \end{aligned}$$

which contradicts f being convex.

□

Lemma 10. *Assume $f(x) = x^2$. Let G be a network chosen in welfare maximizing equilibrium. If G is not connected and the adversary attacks a protected component, then there exists another protected component in G.*

Proof. Assume otherwise. Let (G, Δ^*, A^*) be a welfare maximizing equilibrium. Since the adversary attacks a protected component, there must be unprotected nodes there and the adversary removes some of them. We know that if (G, Δ^*, A^*) is an equilibrium, then $A^*(G, \Delta^*(G))$ does not disconnect the protected nodes. Let P be the protected component and $p = |P|$ be its size, x be the number of unprotected nodes removed, and $u_1, \ldots, u_k$ be the sizes of the remaining, unprotected, components $U_1, \ldots, U_k$ of g, such that $u_1 \geq \ldots \geq u_k$. We will construct a sequence of strategy profiles $(g_i, \Delta_i, A_i)_{0 \leq i \leq l}$ (not necessarily equilibria) such that:

1. $l \geq 1$

2. $(g^0, \Delta^0, A^0) = (G, \Delta^*, A^*)$,

3. (g^l, Δ^l, A^l) is an equilibrium.

4. If $i' < i$, then $W(g^{i'}, \Delta^{i'}, A^{i'}) < W(g^i, \Delta^i, A^i)$.

The points above contradict the assumption that (G, Δ^*, A^*) is a welfare maximizing equilibrium. In each strategy profile (g^i, Δ^i, A^i), Δ^i differs from Δ^* on network g^i only, and A^i differs from A^* on $(g^i, \Delta^i(g^i))$ only. Describing the strategy profiles we will focus on the arguments on which the strategies of the players are different to (G, δ^*, A^*).
Let (g^1, Δ^1, A^1) be defined as follows. Component P is replaced with two components: P_1^1 of size $p_1^1 = p - x'$, and P_2^1 of size $p_2^1 = x'$. Here, $x' = \min Y(x)$, where $Y(x) = \{1 \leq y \leq x : f(y) \geq f(p-y) - f(p-y-1)\}$. The subnetwork of g_1 over P_1^1 is a star and the subnetwork of g_1 over P_2^1 is a clique. Let $\Delta^1(g^1) = \{m\}$, where $m \in P_1^1$ is the centre of the star over P_1^1. Let $A^1(g^1, \Delta^1(g^1)) \in P_2^1$. The construction above is valid as long as x' is well defined, i.e. as long as $Y(x) \neq \varnothing$. This is the case because, by Lemma 8, $x \in Y(x)$.
Given $i \geq 2$, the network (g^i, Δ^i, A^i) is defined on the basis of $(g^{i-1}, \Delta^{i-1}, A^{i-1})$. Each such network has at least two components: P_1^i, of size p_1^i, P_2^i, of size p_2^i. The subnetwork over P_1^i is a star and the subnetwork over the remaining components are cliques. Defence $\Delta_i(g_i) = \{m\}$, where $m \in P_1^i$ is the centre of the star over P_1^i. Attack $A_i(g_i, \Delta_i(g_i)) \in P_2^i$, removes all the component P_2^i. The set of the remaining components is denoted by $\mathcal{U}^i$. The construction ends on minimal i such that for all $U \in \mathcal{U}^i$, $|U| < |P_2^i|$.
Let $t^i = \max\{s \in \mathbb{N} : f(p_2^{i-1}) - f(1) \geq f(p_1^{i-1} + s) - f(p_1^{i-1} + s - 1)\}$. In other words t^i is maximal such that removing all but one node from a component of size $p_2^{i-1} + 1$ is preferred by the adversary to removing a single node from a component of size $p_1^{i-1} + t^i$. Network (g^i, Δ^i, A^i) is obtained from $(g^{i-1}, \Delta^{i-1}, A^{i-1})$ as follows:

1. Pick the largest component U^{i-1} from $\mathcal{U}^{i-1}$. Let $u^{i-1} = |U^{i-1}|$ (note that $u^{i-1} \geq p_2^{i-1}$ as otherwise the algorithm would stop before reaching this point).

2. Move $d^i = \min(t^i, p_2^{i-1} - 1)$ nodes from U^{i-1} to P_1^{i-1} adding them as spokes of the star over P_1^{i-1}, thus obtaining the component P_1^i.

3. Move 1 node from U^{i-1} to P_2^{i-1}, thus obtaining component P_2^i, and form a clique over P_2^i.

Clearly, if $l \geq 2$, the nodes-adversary subgame in the last strategy profile in the sequence, (g^l, Δ^l, A^l), is an equilibrium, as, by the construction, attacking P_2^l is preferred to attacking P_1^l and none of the components in $\mathcal{U}^l$ is larger than $p_2^l - 1$. If the protected node in P_1^l chose no protection, it would be removed by the adversary. Moreover, no node in P_2^l is better off by choosing protection. This is because, by the construction, even if one node protects in P_2^l, the adversary still prefers to attack this component to attacking any other component (note that the components in $\mathcal{U}$ are all strictly smaller than P_2^l).

To see why the nodes-adversary subgame in the last strategy profile in the sequence, (g^l, Δ^l, A^l), is an equilibrium in the case of $l = 1$, notice that no node in P_2^l is better off by deviating and choosing protection. This is because $|P_2^l| = x' \leq x$ and if this is profitable for a node to protect in P_2^l, then it must be profitable to protect for any of the removed nodes neighbouring a protected node in the attacked component P in G (if such a node protects, then the adversary switches his attack and the component of that node in the residual network is of size $\geq x + 1$). This would contradict the assumption that (G, Δ^*, A^*) is an equilibrium.
Now, it is enough to show that $W(G, \Delta^*, A^*) < W(g^l, \Delta^l, A^l)$. To show this we will show, for all $i \in [1, l]$, that

$$W(g^i, \Delta^i, A^i) > W(g^{i-1}, \Delta^{i-1}, A^{i-1}). \tag{42}$$

Clearly for $i = 1$ this is the case, as there are less units of defence used and the number of nodes removed by the adversary is the same or less. For $i \geq 2$ the equation above reduces to

$$f(p_1^{i-1} + d^i) + f(u^{i-1} - d^i - 1) > f(p_1^{i-1}) + f(u^{i-1}). \tag{43}$$

For $f(x) = x^2$ this is equivalent to (substituting, for clarity of presentation, $d \leftarrow d^i$, $p_1 \leftarrow p_1^{i-1}$, $p_2 \leftarrow p_2^{i-1}$, $u \leftarrow u^{i-1}$)

$$(p_1 + d)^2 + (u - d - 1)^2 > p_1^2 + u^2. \tag{44}$$

We show first that $p_1 > u$. To show that it is enough to show that $p - x > u_1$ in $g^0 = G$, as in the subsequent networks p_1^j grows and the sizes of components in $\mathcal{U}^j$ (weakly) decrease, i.e. $p_1 \geq p - x$ and $u \leq u_1$. Since in (G, Δ^*, A^*), the Adversary prefers to attack component P to attacking U_1, so it must be that the residual network with x nodes removed from P has at most the value of the residual network with component U_1 fully removed. Consider network g' obtained from G by re-designing P into a star with centre $m \in P$. Consider the strategy profile (g', Δ', A') where $\Delta'(g') = \{m\}$ and $A'(g', \Delta'(g')) \subseteq U_1$. If the adversary prefers A' to attacking a spoke of P. This profile cannot be an equilibrium, or otherwise G is not optimal for the designer. Thus it must be that the Adversary prefers to attack a spoke of P to attacking a node in U_1, that is

$$p^2 - (p - 1)^2 \geq (u_1)^2 \tag{45}$$

which yields

$$2(p - x) + 2x - 1 \geq (u_1)^2. \tag{46}$$

Notice that it must be that $x \leq u_1$. This is because otherwise the designer would be better off by disconnecting x nodes from P, forming a clique out of them, and changing P into centrally protected star. By Lemma 8, Equation (22), there is an equilibrium in the changed network subgame where the adversary attacks the clique of size x and no node protects in the clique (if a node protected, the adversary would attack the remaining

nodes in the clique). Thus the new network yields a better payoff to the designer, as the loss is the same, but less protection is used.
Since $x \leq u_1$ and $p_1 \geq p - x$, so (46) implies

$$2p_1 + 2u_1 - 1 \geq (u_1)^2. \tag{47}$$

which gives

$$u_1 \leq \sqrt{2p_1} + 1. \tag{48}$$

Since for $p \geq 4$ (and we know that $p \geq 4$ by the structure of the subnetwork of G over P), $p_1 > \sqrt{2p_1} + 1$, so $u_1 < p_1$.
Now, to show (44), we consider two cases separately: (i) $d = p_2 - 1$ and (ii) $d = t^i < p_2 - 1$.
For case (i), we rewrite (44) as

$$(p_1 + p_2 - 1)^2 + (u - p_2)^2 > p_1^2 + u^2 \tag{49}$$

which reduces to

$$(p_2)^2 + (p_2 - 1)^2 > 2(p_2 - 1)(u - p_1) + 2u. \tag{50}$$

Since $(p_2)^2 \geq (p_1)^2 - (p_1 - 1)^2$ (as removing P_2^{i-1} yields better payoff than removing a spoke of the star over P_1^{i-1}), so

$$(p_2)^2 + (p_2 - 1)^2 \geq 2p_1 - 1 + (p_2 - 1)^2. \tag{51}$$

Since $p_1 > u$, so

$$2p_1 - 1 + (p_2 - 1)^2 > 2(p_2 - 1)(u - p_1) + 2u, \tag{52}$$

which implies (44).
For case (ii), Equation (44) can be rewritten as

$$2p_1 d + d^2 > 2u(d + 1) - (d + 1)^2 \tag{53}$$

and further to

$$2(p_1 - 1)d + d^2 > 2ud + d^2 - 2(d^2 - u) - 1. \tag{54}$$

Now we can show that Equation (54) holds. Since

$$(p_2 - 1)^2 - 1 \geq (p_1)^2 - (p_1 - 1)^2 \tag{55}$$

so

$$p_2 \geq \sqrt{2p_1} + 1. \tag{56}$$

Moreover, since

$$(p_2)^2 - 1 < (p_1 + d + 1)^2 - (p_1 + d)^2 \tag{57}$$

(as d is maximal such that attacking P_2^i is preferred to attacking P_1^i) so

$$d > \frac{(p_2)^2}{2} - p_1 - 1 \tag{58}$$

and, by (56),

$$d > \sqrt{2p_1} - \frac{1}{2}. \tag{59}$$

By this

$$d^2 - u > 2p_1 - \sqrt{2p_1} + \frac{1}{4} - u > 0, \tag{60}$$

as $p_1 - 1 \geq u$ and $p_1 + 1 > \sqrt{2p_1}$. Consequently, (54) holds.

By Equation (54), $W(g^i, \Delta^i, A^i) > W(g^{i-1}, \Delta^{i-1}, A^{i-1})$. Thus we have shown that (G, Δ^*, A^*) cannot be an equilibrium, as the designer could choose g^l instead. This completes the proof. □

Lemma 11. *Assume $f(x) = x^2$, and $n \geq 20$. Let (G, Δ, x) be a welfare maximising equilibrium of Γ. If G is not connected, the adversary does not attack a protected component.*

Proof. For a contradiction, suppose that $x(G, \Delta(G)) \in C_1(G)$, where $C_1(G) \cap \Delta \neq \varnothing$. Let $e = |E_{x(G,\Delta(G))}(G|\Delta)|$ denote the number of eliminated nodes, and $|C_1(G)| = y + e$ be the size of the attacked component. By Lemma 6, the attack does not disconnect protected nodes that are connected in G, and so the loss due to attack equals $(e + y)^2 - y^2$.

Pick node i such that $i \in E_{x(G,\Delta(G))}(G|\Delta)$ and $N_i(G) \cap \Delta \neq \varnothing$. That is, node i is eliminated under attack $x(G, \Delta(G))$ and has a protected neighbour. For (G, Δ, x) to be an equilibrium, it must be that $x(G, \Delta(G) \cup \{i\})$ attacks a node in $C_1(G)$ disconnecting protected nodes.[24] It follows that there are at least two protected nodes in $C_1(G)$, i.e. $|C_1(G) \cap \Delta| \geq 2$. Moreover, $e \geq 2$, or otherwise the adversary would strictly prefer attack $x(G, \Delta(G) \cup \{i\})$ to attack $x(G, \Delta(G))$ under $(G, \Delta(G))$. This implies that G cannot feature two isolated nodes; the designer would be strictly better off by connecting them.

By Lemma 11, there exists another component $C_2(G) \neq C_1(G)$ with protected nodes, $C_2(G) \cap \Delta \neq \varnothing$. Without loss of generality, let $C_2(G)$ denote the largest such component. If $C_2(G)$ is not a star, then it can be redesigned as a star, and in this new network there is an equilibrium where only the hub of the star protects, and all other nodes not in $C_2(G)$ choose the same strategy as in the original equilibrium. Moreover, the attack $x(G, \Delta(G))$ is still optimal for the adversary. Since this attains the same gross payoffs with minimal protection spending, let us assume that $C_2(G)$ is a star. Let $z = |C_2(G)|$ denote the size of this component.

We will construct a series of strategy profiles $(G^i, \Delta^i, x^i)_{0 \leq i \leq l}$ such that:

1. $l \geq 1$.

2. $(G^0, \Delta^0, x^0) = (G, \Delta, x)$.

[24] If $x(G, \Delta(G) \cup \{i\}) \notin C_1(G)$, then node i would earn strictly larger payoffs by protecting than the payoffs of its protected neighbour when i does not protect. If $x(G, \Delta(G) \cup \{i\}) \in C_1(G)$ but the attack does not disconnect protected nodes in $C_1(G)$, then node i would earn at least as much as its protected neighbour does when i does not protect. By the tie breaking assumption that a node prefers not to be eliminated, node i would protect.

3. $x^i(G^i, \Delta^i)$ is a best response of the adversary to defence Δ^i in subgame $\Gamma(G^i)$.

4. If $i' < i$, then $W(G^{i'}, \Delta^{i'}, x^{i'}) < W(G^i, \Delta^i, x^i)$.

5. $W(G^l, \Delta^l, x^l) < s^2(n) + (n - s(n) - u(n))$.

In each strategy profile (G^i, Δ^i, x^i), Δ^i differs from Δ on network G^i only, and x^i differs from x on $(G^i, \Delta^i(G^i))$ only. Describing the strategy profiles we will focus on the arguments on which the strategies of the players are different to (G, Δ, x). The points above contradict the assumption that (G, Δ, x) is a welfare maximising equilibrium. If the designer could control protection, then s/he would choose (G^l, Δ^l) over (G^i, Δ^i) for any $i < l$. S/he does not do so because Δ^i is not an equilibrium defence profile. But the network with a star of size $s(n)$, a clique of size $u(n)$ and (possibly) an isolated node achieves strictly higher welfare in equilibrium.
Let (G^1, Δ^1, x^1) be defined as follows. Recall that $|C_1(G)| = y + e$. Take y nodes of $P_1(G)$ are arrange them in a star. Take the remaining e nodes of $C_1(G)$, arrange them in a clique, and link all of the nodes in this clique to the centre of the star of size y. This yields a new component $C_1^1(G^1)$. Let $\Delta^1(G^1) \cap C_1^1(G^1) = \{m\}$, where m is the node that is linked to all other nodes. $x^1(G^1, \Delta^1)$ eliminates the e unprotected nodes attached to m. Note that by construction $x^1(G^1, \Delta^1)$ is a best response of the adversary to defence Δ^1 in sub-game $\Gamma(G^1)$.
Given $i \geq 2$, the network (G^i, Δ^i, x^i) is defined on the basis of $(G^{i-1}, \Delta^{i-1}, x^{i-1})$. Each such network has at least two components: $C_1^i(G^i)$, of size $y + e^i$, and $C_2^i(G^i)$, which is a star of size z^i. Defence of these components is $\Delta^i(G^i) \cap C_i^1(G^i) = \{m\}$, and $\Delta^i(G^i) \cap C_i^2(G^i) = \{h\}$, where h is the centre of $C_2^i(G^i)$. Attack $x^i(G^i, \Delta^i(G^i)) \in C_1^i(G^i)$ removes e^i nodes from $C_1^i(G^i)$.
For $2 \leq i \leq l - 1$, (G^i, Δ^i, x^i) is obtained from $(G^{i-1}, \Delta^{i-1}, x^{i-1})$ as follows:

1. Pick a component $C_j^{i-1}(G^{i-1})$, $j \notin \{1, 2\}$, with $s_j^{i-1} = |C_j^{i-1}(G^{i-1})| \geq 2$. If such a component does not exist, the algorithm stops.

2. Move $t^i = \min\{y + e^{i-1}, s_j^{i-1} - 1\}$ nodes from C_j^{i-1} to C_2^{i-1} adding them as spokes of the star over C_2^{i-1}, thus obtaining component C_2^i.

3. Move 1 node from C_j^{i-1} to C_1^{i-1} adding it to the clique of e^{i-1} nodes to be eliminated, thus obtaining component C_1^i.

4. If, after these changes, there are two isolated nodes, create a component with them.

The construction ends on minimal i such that $|N \setminus \{C_1^i(G) \cup C_2^i(G)\}| \in \{0, 1\}$, i.e. there is at most one node not included in $C_1^i(G)$ or $C_2^i(G)$.
By construction, $x^i(G^i, \Delta^i)$ is a best response of the adversary to defence Δ^i in sub-game $\Gamma(G^i)$. In particular, for every i we have that $f(z^i) - f(z^i - 1) \leq f(y + e^i) - f(y)$. To

see this, first note that it holds for $i = 1$. Then suppose that it holds for $i - 1$:

$$f(z^{i-1})f(z^{i-1} - 1) \leq f(y + e^{i-1}) - f(y). \tag{61}$$

Next note that

$$\begin{aligned} f(z^i) - f(z^i - 1) &= f(z^{i-1} + t^i) - f(z^{i-1} + t^i - 1) \\ &\leq f(z^{i-1} + y + e^{i-1}) - f(z^{i-1} + y + e^{i-1} - 1) \\ &= f(z^{i-1}) - f(z^{i-1} - 1) + 2(y + e^{i-1}) \\ &\leq f(y + e^{i-1}) - f(y) + 2(y + e^{i-1}) \\ &< f(y + e^{i-1}) - f(y) + 2(y + e^{i-1}) + 1 \\ &= f(y + e^{i-1} + 1) - f(y) \\ &= f(y + e^i) - f(y). \end{aligned}$$

Thus, if it holds for $i - 1$, it holds for i. By induction, it holds for all i.
It remains to show that after every application of steps 1-4, the designer is strictly better off. Suppose that $t^i = y + e^{i-1}$. The gain in gross payoffs is bounded below by the case where C_j^{i-1} is of the same size z^{i-1} as C_2^{i-1}, and z^{i-1} is smallest, i.e. $z^{i-1} = t^i + 1$. The gain in this case is of $\{[(t^i + 1) + t^i]^2 - (t^i + 1)^2\} - [(t^i + 1)^2 - 0] = 2(t^i)^2 - 1 > 0$. Suppose next $t^i = s_j^{i-1} - 1$. The gain in gross welfare is of $\left\{[z^{i-1} + (s_j^{i-1} - 1)]^2 - (z^{i-1})^2\right\} - [(s_j^{i-1} - 1)^2 - 0]$, which is greater than zero if and only if $s_j^{i-1} > \frac{1}{2}\frac{2z^{i-1}-1}{z^{i-1}-1}$. Since $\frac{1}{2}\frac{2z^{i-1}-1}{z^{i-1}-1} < 3/2$ and $s_j^{i-1} \geq 2$, the result follows.
Finally, (G^l, Δ^l, x^l) is obtained from $(G^{l-1}, \Delta^{l-1}, x^{l-1})$ as follows. Take the largest number t^l spokes away from C_1^{l-1} and move them as spokes of C_2^{l-1} such that

$$f(z^l) - f(z^l - 1) \leq f(y^l + e^l) - f(y^l), \text{ and} \tag{62}$$

$$f(z^l + 1) - f(z^l) > f(y^l - 1 + e^l) - f(y^l - 1). \tag{63}$$

(62) implies that the adversary's original attack is optimal, whereas (63) implies that if further spokes are moved from C_1^{l-1} to C_2^{l-1} then the adversary would prefer to attack a spoke of C_2^{l-1}. Henceforth, let us denote $z \leftarrow z^l$, $y \leftarrow y^l$, $e \leftarrow e^l$.
Let $s(n) = \lfloor (n+1) - \sqrt{2n} \rfloor$ and $u(n)$ denote the size of the center-protected star and clique, respectively, of the optimal network with at most three components when the adversary attacks an unprotected component. Note that $e < u(n)$ (or otherwise the adversary would attack the clique of size e if it was disconnected), and $z < s(n)$ (or otherwise the adversary would prefer to attack a spoke of the size-z star to eliminating the e nodes attached to the star of size y). We can then consider the following modification to G^l. Out of the y surviving nodes of the attacked component, leave $y - (s(n) - z)$ nodes as unprotected neighbors of the eliminated nodes, and attach the remaining $(s(n) - z)$ nodes as spokes of the other star. The net gain in gross welfare is equal to

$$s^2(n) - (z^2 + y^2) = s^2(n) - z^2 - [(u(n) - e) + (s(n) - z)]^2 = 2\delta_z(z - \delta_e) - \delta_e^2, \tag{64}$$

where $\delta_z \equiv (s(n) - z)$ and $\delta_e \equiv (u(n) - e)$. From (62), we have that $[(u(n) - \delta_e) + (\delta_e + \delta_z)]^2 - (\delta_e + \delta_z)^2 > 2z - 1$, or

$$\delta_e^2 + 2\delta_z\delta_e - [u^2(n) - (2z - 1) + 2u(n)\delta_z] < 0.$$

Since $z < s(n)$, it is easy to see that $u^2(n) - (2z - 1) > 0$: the adversary prefers to eliminate a clique of size $u(n)$ than a spoke of a star of size z. Thus, this polynomial in δ_e has a negative and a positive root. We then have that

$$\delta_e < \bar{\delta}_e \equiv -\delta_z + \sqrt{\delta_z^2 + u^2(n) - (2z - 1) + 2u(n)\delta_z}$$

$\bar{\delta}_e$ is increasing in $u(n)$. Note, however, that $(u(n) - 1)^2 < 2s(n) - 1$, i.e., the adversary prefers to eliminate a spoke of a star of size $s(n)$ to eliminating a clique of $(u(n) - 1)$ nodes. This implies that $u(n) < 1 + \sqrt{2s(n) - 1}$, so that

$$\delta_e < -\delta_z + \sqrt{\delta_z^2 + 4\delta_z + 1 + 2(1 + \delta_z)\sqrt{2s(n) - 1}}.$$

Combining this result with (64) yields

$$\begin{aligned} s^2(n) - (z^2 + y^2) &> \delta_z(2z - 4) - 2(1 + \delta_z)\sqrt{2s(n) - 1} - 1 \\ &\geq \delta_z(2z - 4) - 2(1 + \delta_z)u(n) - 1 \\ &= 2\delta_z(\delta_z + [z - (u(n) + 2)]) - 2\delta_z^2 - 2u(n) - 1 \qquad (65) \end{aligned}$$

where the second inequality uses $\sqrt{2s(n) - 1} \leq u(n)$. Recall that $\delta_z \geq 1$. The proof is completed with with following two steps. *First*, let us show that $z \geq u(n) + 2$, so that the right-hand side of (65) is minimized at $\delta_z = 1$. Suppose, for a contradiction, that $z \leq u(n) + 1$. It is straightforward to verify that $\frac{u(n)+1}{n} \leq 0.4$ for any $n \geq 15$. Thus, for $n \geq 15$ we have that $\frac{z}{n} \leq 0.4$. On the other hand, note that $y < z/2$. To see this, note that (63) can be written as $y < \frac{z - \frac{1}{2}(e-1)^2}{e} < z/2$, where the second inequality uses $e \geq 2$. Moreover, $x < z$ (or otherwise the adversary would attack a disconnected clique of size x rather than a spoke a star of size z). Therefore, $z/n = z/(x + y + z) > z/(z + z/2 + z) = 0.4$, a contradiction. *Second*, substituting $\delta_z = 1$ in (65), we obtain $s^2(n) - (z^2 + y^2) > 2s(n) - 4u(n) - 7$, where the right-hand side is positive for $n \geq 20$. □

Lemma 12. *Assume $f(y) = y^2$. If G is not connected, $x(G, \Delta) \in C_j$ where $C_j \cap \Delta = \varnothing$, and $\Delta \neq \varnothing$, then:*

(a) $|\Delta| = 1$, i.e. there is only one protected node.

(b) There are at most two unprotected components.

(c) If there are two unprotected components, one of them is of size 1.

Proof. Let $P_1, \ldots, P_l$ denote components with at least one protected node, and $C_1, \ldots, C_r$ denote unprotected components. Component labels are such that $|P_1| \geq |P_2| \geq, \ldots, \geq |P_l|$ and $|C_1| \geq |C_2| \geq, \ldots, \geq |C_r|$.

By Lemma 9, there is no fully protected component. Therefore, if P_i is not a star, the designer can re-design it as a star, and the profile where only the centre protects is an equilibrium profile. Since this achieves the same connectivity with minimal protection, the designer is not worse off. We can then assume that P_i is a star for all i. Let $|P_1| = s$ and $|C_1| = u$ denote the sizes of the largest star and the largest unprotected component, respectively. For the adversary to attack C_1 in equilibrium, it must be that $u^2 \geq 2s - 1$ if $c > u$, or $u^2 - 1 \geq 2s - 1$ if $c \leq u$. Since $u^2 - 1 \geq 2s - 1$ implies $u^2 \geq 2s - 1$, let us assume the more restrictive case where u is the smallest integer such that $u^2 - 1 \geq 2s - 1$.

Specifically, let u be the smallest integer such that $2s - 1 \leq u^2 - 1 < 2s + 1$. If $n = s + u$ or $n = s + u + 1$, then (a), (b) and (c) hold. We will show that if $n > s + u + 1$, i.e. there are at least two other nodes in the network, then G cannot be optimal.

Note that if the designer adds an additional node in $N/(P_1 \cup C_1)$ to C_1, she could attach a maximum of $\mathcal{S}(u)$ additional spokes to P_1 such that the adversary would still strictly prefer to attack the unprotected component, where

$$\mathcal{S}(u) = \begin{cases} u & \text{if } u \text{ is even} \\ u + 1 & \text{if } u \text{ is odd} \end{cases} . \tag{66}$$

Further note that $u \geq 2$: if $u = 1$ then eliminating the spoke of a star will always be preferred. It follows that, if the size of the clique is increased by 1, then the size of the center-protected star can be increased by at least 1 and the adversary's original attack will remain optimal.

Recall that $n > s + u + 1$. If all other components in G are of size 1, then consider the following modifications to G yielding a network G'. Create a clique including all nodes originally in C_1 and one additional node who was isolated in G, and add a spoke to P_1. By (66), the adversary attacks the clique of G' of size $(u + 1)$. Since $s \geq 2$, $(s + 1)^2 - (s^2 + 2) > 0$, and gross welfare is strictly higher under G' than under G. Since protection spending does not change, G' achieves strictly higher welfare than G, a contradiction.

Suppose finally that there is a component K of size k, $1 < k \leq s$. Then gain in welfare by re-allocating nodes in K to components C_1 and P_1 is bounded below by the case where $k = 2$: only one spoke can be added to P_1 after adding one node to C_1. Consider then network G'', obtained from G by creating a clique including all nodes originally in C_1 and one of the nodes of K, and adding the remaining node of K as a spoke of P_1. By (66), the adversary attacks the clique of G'' of size $(u + 1)$. Gross gain in welfare is equal to $(s + 1)^2 - (s^2 + 4) > 0$. Since protection spending remains constant or decreases, G'' achieves strictly higher welfare than G, a contradiction. $\square$

It follows from Lemma 12 that if the adversary attacks an unprotected component, then G consists of a star of size $s(n)$ and a component of size $u(n)$, and $\Delta = \{m\}$, where $\{m\}$ is the centre of the star.

Proof of Proposition 2. By Proposition 4 and Lemmas 11-12, three architectures and defence profiles are possible under welfare minimizing equilibria: a connected network where all nodes protect, a disconnected network as described in point (2) of the proposition, and a disconnected network as described in point (3) of the proposition. Comparing welfare attained in (1)-(3) yields the thresholds $c_D(n)$ and $c_U(n)$. In particular, the network and defence profile achieving highest welfare is (1) if $0 < c \leq \min\{c_D(n), c_U(n)\}$, (2) if $c_D(n) < c < c_U(n)$, and (3) if $c \geq \max\{c_D(n), c_U(n)\}$. However, defence profile in (2) is equilibrium defence profile if and only if $c \leq \frac{f(s(n))}{s(n)} = s(n)$ (so that the centre protects). Considering this yields the desired result. □

D.2 Welfare minimizing equilibria

As in the proof of existence (Lemma 1), given network G and costs c, we will use $\mathcal{A}(G|c)$ to denote the set of components in G were it is not individually rational to protect under any attack strategy. That is, for any $C \in \mathcal{A}(G|c)$, $f(|C|)/|C| < c$. The following fact will be used to prove some of the results.

Fact 5. *Let (G^*, Δ, x) be a welfare minimizing equilibrium. If G^* is disconnected, then $\mathcal{A}(G^*|c) \neq \varnothing$.*

Proof. Suppose there is no such component. Let $\Delta^* = \Delta(G^*)$ and $x^* = x(G^*, \cdot)$. By the construction used in proof of Lemma 1, there exists an equilibrium of $\Gamma(G^*)$, such that all nodes protect. Since (Δ^*, x^*) is welfare minimizing, this equilibrium is not worse for **D**. Let G' be a star over all nodes from N. Since $f(n-1)/(n-1) \geq c$, in any equilibrium (Δ', x') of $\Gamma(G')$, $\Delta' = N$ and no node is infected by $x'(\Delta')$. Moreover, by convexity of f, **D** is strictly better off than under G^*, a contradiction with our assumptions. Thus it must be that there exists $X \in \mathcal{C}(G^*)$ such that $f(|X|)/|X| < c$. □

We start by showing that, in a welfare minimizing equilibrium, there cannot be a fully protected component.

Lemma 13. *Let (G^*, Δ, x) be a welfare minimizing equilibrium. If G^* is disconnected, then there is no fully protected component.*

Proof. Assume, to the contrary, that there exists $C \in \mathcal{C}(G^*)$ such that $C \subseteq \Delta$. Let $\Delta^* = \Delta(G^*)$ and $x^* = x(G^*, \cdot)$. As in the proof of existence (Lemma 1), given network G and costs c, we will use $\mathcal{A}(G|c)$ to denote the set of components in G were it is not individually rational to protect under any attack strategy. By Fact 5, $\mathcal{A}(G^*|c) \neq \varnothing$. Moreover, in any equilibrium of $\Gamma(G^*)$ and for any $C \in \mathcal{A}(G^*|c)$, no node protects in C under $\tilde{\Delta}$.

Let G' be a network defined as follows. The sets of components of G' and G^* are the same, $\mathcal{C}(G') = \mathcal{C}(G^*)$, and for each $X \in \mathcal{C}(G')$, $G'[X]$ is a star. We will show that either G' yields higher payoff than G^* under welfare minimizing equilibria to **D**, or there exists G'' that does. This will contradict our assumptions and complete the proof. The proof goes by steps.
Let Δ^{ex} be an equilibrium defence of $\Gamma(G^*)$, constructed as in proof of equilibrium existence (Lemma 1).

Step 1. In any equilibrium (Δ', x') of $\Gamma(G')$, for any $X \in \mathcal{C}(g) \setminus \mathcal{A}(G'|c)$, $X \cap \Delta' \neq \varnothing$. For assume otherwise. Then there exists $X \in \mathcal{C}(G) \setminus \mathcal{A}(G'|c)$ such that $X \cap \Delta' = \varnothing$. Clearly **A** attacks one such X of maximal size. It must be that $x'(\Delta' \cup \{i\}) \in X \setminus \{i\}$, where i is the centre of $G'[X]$, for otherwise it would be profitable for the centre of X to protect (because $f(|X|)/|X| \geq c$ and every node prefers outcomes where it stays uninfected). Similarly, it must be that $f(|X|-1)/(|X|-1) < c$, as otherwise deviation to protection would be profitable to i. But then X is a component in $\mathcal{C}(G') \setminus \mathcal{A}(G'|c)$ of minimal size (if it had one node less, it would be in $\mathcal{A}(G'|c)$). Notice also that it must be that $X \subseteq \Delta^{ex}$, for otherwise it would mean that **A** prefers attacking a component in $\mathcal{A}(G^*|c)$ than removing a node in X (removing a node is the smallest possible damage that **A** can cause to X when X is not fully protected). Furthermore, since X is of minimal size to be in $\mathcal{C}(G') \setminus \mathcal{A}(G'|c)$, it must be that *all* nodes in $N \setminus \mathcal{A}(G^*|c)$ are protected under Δ^{ex}. Clearly, this means that X is the unique component in $\mathcal{C}(G') \setminus \mathcal{A}(G'|c)$, or otherwise a network G'' that merges all components in $\mathcal{C}(G') \setminus \mathcal{A}(G'|c)$ into a single star attains (by convexity of f) strictly higher payoffs than G to **D** in any equilibrium, for in any equilibrium of $\Gamma(G'')$ all nodes in the star protect.
Note that there must exist at least two components in $\mathcal{A}(G'|c)$, for otherwise the unique unprotected components, Z, yields zero payoffs to **D**, and **D** is strictly better off by choosing a star network where all nodes protect in any equilibrium. The payoffs of nodes in X increase, as well as the payoffs of nodes in Z.
Consider then the following two modifications to G'. Network $\hat{G}$ is obtained by attaching component Y to X and forming a star component, where $x(G', \Delta^{ex}) \notin Y$ (such an unattacked component exists because of the argument in the preceding paragraph). In any equilibrium of $\hat{G}$ all nodes in the star component protect. Let $|X| = |\tilde{X}| + |Y|$. For this not to be profitable, it must be that $c|Y| \geq f(|\tilde{X}|+2|Y|) - f(|\tilde{X}|+|Y|) - f(|Y|)$. Network $\tilde{G}$ is formed by creating a star out of X and detaching $|Y|$ spokes to form a copy of Y. Since X was of minimal size for it to be individually rational to protect, clearly no node protects in any equilibrium of $\Gamma(\tilde{G})$. Two cases are possible. If **A** does not attack $\tilde{X}$, then this is not a profitable modification to **D** iff $(|\tilde{X}|+|Y|)c \leq f(|\tilde{X}|+|Y|) - f(|\tilde{X}|) - f(|Y|)$. Combining with the condition that states that $\hat{G}$ is not profitable yields $f(|\tilde{X}|+2|Y|) - f(|\tilde{X}|+|Y|) \leq f(|\tilde{X}|+|Y|) - f(|\tilde{X}|)$, a contradiction with f being convex. If **A** attacks $\tilde{X}$, then this is not a profitable modification to **D** iff $(|\tilde{X}|+|Y|)c \leq f(|\tilde{X}|+|Y|) - f(|Y|) - f(|Z|)$,

where Z is the originally attacked component. Since it is not individually rational to protect in $\tilde{X}$, $\tilde{X}c > f(|\tilde{X}|)$, so that $f(|\tilde{X}|) + |Y|c < f(|\tilde{X}| + |Y|) - f(|Y|) - f(|Z|)$, or $|Y|c < f(|\tilde{X}| + |Y|) - f(|\tilde{X}|) - f(|Y|) - f(|Z|)$. Combining with condition for $\hat{G}$ not to be profitable yields $f(|\tilde{X}| + 2|Y|) - f(|\tilde{X}| + |Y|) < f(|\tilde{X}| + |Y|) - f(|\tilde{X}|) - f(|Z|) < f(|\tilde{X}| + |Y|) - f(|\tilde{X}|)$, a contradiction with f being convex. Therefore, in any equilibrium either $\hat{G}$ or $\tilde{G}$ attain strictly higher welfare than G' does under Δ^{ex}, a contradiction with G' being optimal under welfare minimizing equilibrium.

Step 2. In any equilibrium (Δ', x') of $\Gamma(G')$, for any $X \in \mathcal{C}(G) \setminus \mathcal{A}(G'|c)$, either $X \subseteq \Delta'$ or $X \cap \Delta = \{i\}$, where i is the centre of $G'[X]$.

Notice first that there exists a component $X \in \mathcal{C}(G^*) \setminus \mathcal{A}(G^*|c)$ such that $X \subseteq \Delta'$. To see this take any $X \in \mathcal{C}(G^*) \setminus \mathcal{A}(G^*|c)$ such that $X \subseteq \Delta^*$ (as we assumed, such a component exists). It must be that removing a single node from X is preferred by **A** to attacking a largest component in $\mathcal{A}(G^*|c)$, for such an attack is available to **A** under Δ^* on G^* and yet all nodes in X protect under Δ^*. Since this is the case in G' as well, so either none or all nodes protect in X under Δ'. We ruled out the former in Step 2. Hence it must be that $X \subseteq \Delta'$.

Now, suppose, to the contrary of the statement in Step 3, that there exists $X \in \mathcal{C}(g) \setminus \mathcal{A}(G'|c)$ such that neither $X \subseteq \Delta'$ nor $X \cap \Delta' = \{i\}$. Since $X \cap \Delta' = \varnothing$ is ruled out by Step 2, it must be that $2 \leq |X \cap \Delta'| < |N|$.

As the first case, suppose that $X \setminus \{i\} \subseteq \Delta'$, where i is the centre of $G'[X]$. In other words, all spokes of $G[X]$ protect and its centre does not. Let $|Y|$ be the the component in $\mathcal{A}(G|c)$ attacked by $x'(\Delta')$ and Z be a largest component fully protected under Δ'. Since attacking Y is preferred to attacking the centre of X, so

$$f(|Y|) \geq f(|X|) - (|X| - 1)f(1). \tag{67}$$

On the other hand, since Z is fully protected, so

$$f(|Z|) - f(|Z| - 1) \geq f(|Y|). \tag{68}$$

Thus

$$f(|Z|) - f(|Z| - 1) \geq f(|X|) - (|X| - 1)f(1), \tag{69}$$

and, by convexity of f,

$$f(|Z| + 1) - f(|Z|) > f(|X|) - (|X| - 1)f(1), \tag{70}$$

so

$$f(|Z| + 1) + (|Z| - 1)f(1) > f(|X|) + f(|Z|). \tag{71}$$

Consider network G'' obtained from G' by disconnecting $G'[X]$ and attaching one node from X to $G[Z]$ as a spoke. In any equilibrium on G'', the extended Z' fully protects and none of the remaining nodes from X protect. Moreover, any welfare minimizing equilibrium on G'' translates to a welfare minimizing equilibrium on G', where the same nodes protect apart from those in $X \cup Z$. The same operation may be applied to get rid of all components which have spokes-only-protect equilibria on G'. By (71), the value of the network is strictly increasing between G' and G''. Now, all fully protected components under Δ^* are also fully protected under Δ^{ex} and in any equilibrium on G' and G''. All the other protected components are replaced with centrally protected stars in G'' or are profitably merged with fully protected components (and during this merging the number of protected nodes between Δ^{ex} and any equilibrium on G'' does not increase). Hence G'' is strictly better to G^* under welfare minimizing equilibria.

Steps 1-2 establish that, in any equilibrium of $\Gamma(G')$, there is a unique fully protected component X, and all other components in $\mathcal{C}(G) \setminus \mathcal{A}(G'|c)$ are centre protected stars. Note that under equilibrium defence Δ^{ex} in $\Gamma(G')$, $\mathbf{D}$ attains the same gross welfare with at least the same protection spending. For, by construction, $X \subseteq \Delta^{ex}$ and all other components in $\mathcal{C}(G) \setminus \mathcal{A}(G'|c)$ have at least one unit of protection. Therefore, we have that $U^{\mathbf{D}}(G^*, \Delta^*, x^*(\Delta^*)) \leq U^{\mathbf{D}}(G^*, \Delta^{ex}, x^*(\Delta^{ex})) \leq U^{\mathbf{D}}(G', \Delta', x'(\Delta'))$.

Note that there must exist at least two components in G', for otherwise attaching the attacked nodes as spokes of X would make $\mathbf{D}$ strictly better off. Thus, there exists an unattacked component $Y \neq X$. The proof is completed by considering two different modifications, yielding to networks $\hat{G}$ and $\tilde{G}$. Network $\hat{G}$ is obtained by attaching component Y to X and forming a star component. Network $\tilde{G}$ is formed by creating a star out of X and detaching $|Y|$ spokes to form a copy of Y. Following analogous steps as in the proof of Lemma 9 shows that under welfare minimizing equilibria either $\hat{G}$ or $\tilde{G}$ make $\mathbf{D}$ strictly better off than under G', contradicting G^* being optimal.

□

Proof of Proposition 5. For a contradiction, let (G^*, Δ, x) be a welfare minimizing equilibrium where, for $P \in \mathcal{C}(G)$, $x(G, \Delta(G)) \in P$ and $P \cap \Delta(G) \neq \varnothing$. Let $\Delta^* = \Delta(G^*)$ and $x^* = x(G^*, \cdot)$. Clearly, $P \not\subseteq \Delta^*$. For otherwise the adversary does not eliminate a single node, and it must be that $\Delta^* = N$. But then a connected star network attains strictly higher payoffs to $\mathbf{D}$ in any equilibrium, as all nodes protect as well but (due to convexity of f) gross payoffs are higher. Thus, $P \not\subseteq \Delta^*$ and $\mathbf{A}$ eliminates at least one node in P. Moreover, by Fact 5, $\mathcal{A}(G^*, c) \neq \varnothing$.

Let X denote the set of eliminated nodes in equilibrium (G^*, Δ, x), i.e. $X = E_x(G^*|\Delta)$. Moreover, let $\Delta^i = \Delta^* \cup \{i\}$, $x^i = x(G^*, \Delta^i)$ be equilibrium response of $\mathbf{A}$ to G^* and defence Δ^i, and let $X^i = E_{x^i}(G^*|\Delta^i)$ be the set of nodes eliminated by x^i.

Step 1. For any $i \in \partial_G(\Delta^*) \cap X$, $|\mathcal{C}(G[P] - X^i)| \geq 2$. That is, if an eliminated node with a protected neighbor protects, the best response of $\mathbf{A}$ results in the residual network over P having at least two components. Suppose, to the contrary, that component $G[P]$ does not get disconnected by x^i. It must be that $X^i \subseteq P$ and $X^i \not\subseteq X$. If $X^i \not\subseteq P$, then node i would prefer to protect, since $f(|P|)/|P| - c > f(|P| - |X|)/(|P| - |X|) - c \geq 0$, where the first inequality if by f increasing and convex and $f(0) = 0$, and the second inequality by the fact that there are protected nodes in P in equilibrium (G^*, Δ, x). If $X^i \subseteq X$, then by protecting node i gets a payoff of at least $f(|P| - |X| + 1)/(|P| - |X| + 1) - c > f(|P| - |X|)/(|P| - |X|) - c \geq 0$. Moreover, $|X^i| \leq |X|$ (as otherwise x^i would be a better response to (G, Δ^*) than x) and so payoff to i in $G[P] - X^i$, when it protects, is $f(|P| - |X^i|)/(|P| - |X^i|) - c \geq f(|P| - |X|)/(|P| - |X|) - c \geq 0$, so i is better off. Thus $G[P]$ must get disconnected by x^i, i.e. $|\mathcal{C}(G[P] - X^i)| \geq 2$.

Step 2. For any $\tilde{\Delta}$ such that $|\tilde{\Delta} \cap P| \leq 1$, $x(G^*, \tilde{\Delta}) \notin C$ for any $C \in \mathcal{A}(G^*, c)$. If $|\tilde{\Delta} \cap P| = 0$, then $P \notin \mathcal{A}(G^*, c)$ implies that A must strictly prefer eliminating P to eliminating $C \in \mathcal{A}(G^*, c)$. Suppose then $|\tilde{\Delta} \cap P| = 1$. As in the previous paragraph, let $X = E_x(G^*|\Delta)$ denote the set of eliminated nodes in equilibrium (G^*, Δ, x), and $X^i = E_{x^i}(G^*|\Delta^i)$ the set of nodes eliminated if node $i \in X$ protects. Let U denote the largest component in $\mathcal{A}(G^*, c)$. Note that eliminating a component $C \in \mathcal{A}(G^*, c)$ is available to A under (G^*, Δ^*). $x^* \notin \mathcal{A}(G^*, c)$ implies that $f(|U|) \leq f(|P|) - f(|P| - |X|)$ (the damage caused by attacking X is at least as large as that of attacking U). $X^i \subseteq P$ and $X^i \not\subseteq X$ implies that, for any $i \in X$, $f(|U|) \leq f(|P|) - f(|X \cup \partial_G(X)|) - \sum_{C \in \mathcal{C}(G[P] - X^i) \setminus (X \cup \partial_G(X))} f(|C|)$ (the damage caused by attacking X^i is at least as large as that of attacking U). Moreover, $|\mathcal{C}(G[P] - X^i)| \geq 2$ (by Step 1) implies that, for any $i \in X$, $C \in \mathcal{C}(G[P] - X^i) \setminus (X \cup \partial_G(X)) \neq \varnothing$ and thus the last term $\sum_{C \in \mathcal{C}(G[P] - X^i) \setminus (X \cup \partial_G(X))} f(|C|) > 0$. There are two cases to consider. Case (i): $\tilde{\Delta} \cap P \subseteq X \cup \partial_G(X)$, i.e. the protected node in P is a node in X or has a neighbor in X. Pick any $i \in X$, and consider an attack on node $j \in X^i$. Damage caused by A bounded below by case where $\tilde{\Delta} \cap P = \partial_G(X)$, in which case it is of $f(|P|) - f(|X \cup \partial_G(X)|) > f(|P|) - f(|X \cup \partial_G(X)|) - \sum_{C \in \mathcal{C}(G[P] - X^i) \setminus (X \cup \partial_G(X))} f(|C|) \geq f(|U|)$. Hence A strictly prefers an attack on $j \in X^i$ to an attack on U, and therefore, $x(G^*, \tilde{\Delta}) \notin C$ for any $C \in \mathcal{A}(G^*, c)$. Case (ii): $\tilde{\Delta} \cap P \not\subseteq X \cup \partial_G(X)$. Damage caused by an attack on any $i \in X$ is of at least $f(|P|) - f(|P| - |X| - 1) > f(|P|) - f(|P| - |X|) \geq f(|U|)$. Hence A strictly prefers an attack on $i \in X$ to an attack on U, and therefore, $x(G^*, \tilde{\Delta}) \notin C$ for any $C \in \mathcal{A}(G^*, c)$.

Step 3. G^* is not optimal. The proof is finalized with the following arguments. Let Δ^{ex} be an equilibrium defence of $\Gamma(G^*)$, constructed as in proof of equilibrium existence (Lemma 1). Since $\mathcal{A}(G^*, c) \neq \varnothing$, $x(G^*, \Delta^{ex}) \in C$, where $C \in \mathcal{A}(G^*, c)$, i.e. in equilibrium A eliminates a component in $\mathcal{A}(G^*, c)$. Note that, since (Δ, x) is welfare minimizing on G^*, $U^{\mathsf{D}}(G^*, \Delta^*, x^*(\Delta^*)) \leq U^{\mathsf{D}}(G^*, \Delta^{ex}, x^*(\Delta^{ex}))$. Two cases must be considered.

Case (a). There exists $Z \in \mathcal{C}(G^*)$ such that $Z \subseteq \Delta^{ex}$. Then note that there must exist at least two components in G^*, for otherwise a connected star attains strictly higher payoffs to D in any equilibrium. Thus, there exists an unattacked component $Z' \neq Z$. Consider then two different modifications, yielding to networks $\hat{G}$ and $\tilde{G}$. Network $\hat{G}$ is obtained by attaching component Z' to Z and forming a star component. Network $\tilde{G}$ is formed by creating a star out of Z and detaching $|Z'|$ spokes to form a copy of Z'. Following analogous steps as in the proof of Lemma 9 shows that under welfare minimizing equilibria either $\hat{G}$ or $\tilde{G}$ make **D** strictly better off than under G', contradicting G^* being optimal.

Case (b). There is no fully protected component under Δ^{ex}. By construction of Δ^{ex}, there is at least one protected node in every component $C \in \mathcal{G}^* \setminus \mathcal{A}(G, c)$. Moreover, since $x(G^*, \Delta^{ex}) \in C$ for some $C \in \mathcal{A}(G^*, c)$, by Step 2 above it must be that $|\Delta^{ex} \cap P| \geq 2$. Consider then network G', obtained from G^* as follows. The sets of components of G' and G^* are the same, $\mathcal{C}(G') = \mathcal{C}(G^*)$, and for each $X \in \mathcal{C}(G')$, $G'[X]$ is a star. Consider defence profile Δ' where, for each $C \in \mathcal{C}'(G') \setminus \mathcal{A}(G', c)$, $\Delta' \cap C = \{i\}$ where i is the centre of C. This defence profile is the unique equilibrium defence of $\Gamma(G')$. For any node in eliminated component $U \in \mathcal{A}(G', c)$, not to protect is a strictly dominant strategy. Consider next any of the stars $P' \notin \mathcal{A}(G', c)$. Even if all spokes protect, **A** prefers to attack the centre of the star to attacking a component any $C \in \mathcal{A}(G', c)$. To see this, note that convexity of f and $\frac{f(|U|)}{|U|} < c \leq \frac{f(|P'|)}{|P'|}$ imply that $|U| \leq |P'| - 1$, and so $f(|U|) \leq f(|P'| - 1)$. Next node that, by Property 1, $f(|P'| - 1) < f(|P'|) - (|P'| - 1)f(1)$. Combining, we have that $f(|U|) < f(|P'|) - (|P'| - 1)f(1)$, which implies that **A** would strictly prefer to attack the centre of the star of size $|P'|$ to attacking the unprotected component of size $|U|$ even if all spokes of P' protect. Hence, the centre of P' protects in any equilibrium. If the centre of the star protects, then spokes do not protect, as **A** prefers attack of U to eliminating a single node of P' that does not disconnect P'. The proof is finalized by noting that $(G', \Delta', x(G', \Delta'))$ achieves the same gross welfare than (G, Δ^{ex}, x^*) but with strictly smaller protection spending, since $|\Delta' \cap P| = 1 < 2 \leq |\Delta^{ex} \cap P|$. Thus, $U^{\mathbf{A}}(G', \Delta', x(G', \Delta')) > U^{\mathbf{A}}(G^*, \Delta^{ex}, x^*) \geq U^{\mathbf{A}}(G^*, \Delta^*, x^*)$, and so G^* cannot be optimal. □

We can now use the results derived for general network value function to obtain Proposition 3 and Corollary 1.

Proof of Lemma 4. For $n \geq 4$, $c_U(n) \leq n - 1$. Therefore, $c \leq c_U(n)$ implies $c \leq n - 1$, and **D** can choose the star network G where $\Delta(G) = N$ in any equilibrium (Δ, x) of $\Gamma(G)$. □

Proof of Proposition 3. Let (G, Δ, x) be a welfare minimizing equilibrium. By Lemma 13 and Proposition 5, if G is disconnected but $\Delta(G) \neq \varnothing$, then there is no fully protected component and **A** attacks an unprotected component. By Lemma 12, in this case G has only one protected node. Clearly, the protected component must be a star of maximal size, $\hat{s}(n)$, such that **A** strictly prefers to attack the unprotected component.
Thus, three architectures can be optimal. If G is connected then $\Delta(G) = N$. If G is disconnected but $\Delta(G) \neq \varnothing$, then G features a star of size $\hat{s}(n)$ and an unprotected

component of size $u(n)$. If G is disconnected and $\Delta(G) = \varnothing$, then G is the optimal unprotected network. Comparing payoffs yields thresholds $\hat{c}_D(n)$ and $c_U(n)$.
To see that defence profiles are equilibrium profiles, first note that if $c \leq \min\{\hat{c}_D(n)c_U(n)\}$ then by Lemma 4 there exists G such that $\Delta(G) = N$ in any equilibrium of $\Gamma(G)$. Furthermore, if $c > \max\{c_U(n), \hat{s}(n)\}$ then $c > f(\lfloor n/2 \rfloor) / (\lfloor n/2 \rfloor)$, and so if G is the optimal unprotected network then $\Delta(G) = \varnothing$ in any equilibrium of $\Gamma(G)$. Finally, if $u(n) < c_D(n)$ implies that if $c_D(n) < c \leq \hat{s}(n)$ then in any equilibrium of $\Gamma(G)$ only the centre of the star protects and **A** eliminates unprotected component of size $u(n)$. □

Proof of Corollary 1. Let (G, Δ, x) be a welfare minimizing equilibrium. We first show that if G is disconnected then $\Delta(G) = \varnothing$. By Lemma 13 and Proposition 5, if G is disconnected but $\Delta(G) \neq \varnothing$, then there is no fully protected component and **A** attacks an unprotected component. Let $U \in \mathcal{C}(G)$ denote the attacked component, and $P \in \mathcal{C}(G)$ a partially protected component. If $|P| > |U|$, then $f(y) > 2f(y-1)$ implies that **A** must strictly prefer an attack on an unprotected node in P than eliminating U. Hence it must be that $|P| \leq |U|$. If $|P| \leq |U|$ and some nodes protect in P, then there exists an equilibrium (Δ', x') of $\Gamma(G)$ where $P \cup U \subseteq \Delta'(G)$. Since (Δ, x) is welfare minimizing on G, **D** cannot be worse off. But then consider network G'' where P and U are merged into a star. All nodes in $P \cup U$ protect in any equilibrium of $\Gamma(G'')$ and, by convexity of f, **D** is strictly better off. Therefore, if G is disconnected then $\Delta(G) = \varnothing$.
We thus have that G is connected and $\Delta(G) = N$, or G is the optimal unprotected network and $\Delta(G) = \varnothing$. Comparing payoffs indicates that **D** prefers full protection if $c \leq \frac{2^n - 2\lfloor \frac{n}{2} \rfloor - n \bmod 2}{n}$, and no protection otherwise. However, for $c > \underline{c}(n)$ every network has a no protection equilibrium. Since $\underline{c}(n) \leq \frac{2^n - 2\lfloor \frac{n}{2} \rfloor - n \bmod 2}{n}$ (with strict inequality if $n \geq 3$), **D** chooses connected G such that $\Delta(G) = N$ in any equilibrium of $\Gamma(G)$ if $c \leq \underline{c}(n)$, and the optimal unprotected network otherwise. □

D.3 Securing full protection through network design

Proof of Proposition 6. Consider first (1). For a contradiction, suppose there is no k-critical node with $\frac{f(n-k)}{n-k} > c$. Consider the strategy profile (Δ, x) in which no node protects. Pick any node $i \in V$. If i does not protect, it gets a payoff 0. If i protects, it gets a payoff of $U^i(G, \Delta, x(\Delta)) = \frac{f(n-k_i)}{n-k_i} - c$, where k_i is the size of the largest component in $G - \{i\}$. Since there is no k-critical node with $\frac{f(n-k)}{n-k} \geq c$ so $U^i(G, \Delta, x(\Delta)) < 0$. Therefore, the profile in which no node protects is an equilibrium.
Now consider (2). Let (Δ, x) be an equilibrium of $\Gamma(G)$ where not all nodes protect. Let $i \in E_{x(\Delta)}(G|\Delta)$. Since $\Delta \subsetneq N$, $E_{x(\Delta)}(G|\Delta) \neq \varnothing$. It cannot be that i is a k-critical node with $\frac{f(n-k)}{n-k} > c$, as i would prefer to protect. Suppose that i is not k-critical with $\frac{f(n-k)}{n-k} > c$. Then i is connected to a k-critical node j with $\frac{f(n-k)}{n-k} > c$. If i deviates to protection, its payoff will be $U^i(G, \Delta \cup \{j\}, x(\Delta \cup \{j\})) \geq \frac{f(n-k_j)}{n-k_j} - c$, where k_j is the size

of the largest component in $G - \{j\}$. Since $U^i(G, \Delta \cup \{j\}, x(\Delta \cup \{j\})) > 0$, so i is better off by deviating, which contradicts the assumption that (Δ, x) is an equilibrium. $\square$

E Proofs for random attack

E.1 First best

Proof of Proposition 7. Let (G, Δ) be a first best protected network. Three cases are possible.

Case (i). $\Delta = N$ Clearly in this case G must be a connected network.

Case (ii). $\varnothing \subsetneq \Delta \subsetneq N$ Then (G, Δ) must be a centre-protected star. We prove this in three steps.

Step 1. G is connected. For a contradiction, suppose G is not connected. Since $\Delta \neq \varnothing$, there exists $C(G) \in \mathcal{C}(g)$ with a node $i \in C(G) \cap \Delta$. Consider the following modification to G, which results in network G'. For every node $j \notin C(G)$, delete all its links and create a link between j and i. Protection spending remains the same, and gross expected payoffs from connectivity strictly increase. Hence G cannot be optimal.

Step 2. If $k, j \notin \Delta$, then $kj \notin G$. That is, there are no links between unprotected nodes. Suppose the contrary. If k and j are not leaf nodes, then consider network G' which is identical to G except that $kj \notin G'$. Gross expected payoffs are strictly greater under G' than under G. Suppose that k is a leaf. Let $i \in \Delta$ be a protected node. Consider network G'' which is identical to G except that $kj \notin G''$ and $ki \in G''$. Gross expected payoffs are strictly greater under G'' than under G.

Step 3. (G, Δ) is a centre-protected star. Let s denote the number of protected nodes. The designer's payoffs are equal to

$$\frac{n-s}{n} f(n-1) + \frac{s}{n} f(n) - sc = f(n-1) + s\left[\frac{f(n) - f(n-1)}{n} - c\right].$$

Note that is must be that $\frac{f(n)-f(n-1)}{n} < c$, or otherwise it would be optimal to protect all nodes, a contradiction. The payoff of the designer is therefore maximised at $s = 1$. That is, a single node is protected, and thus (G, Δ) is a centre-protected star.

Case (iii). $\Delta = \varnothing$ Since there are at most n components in the network, the designer solves

$$\arg\max_{\boldsymbol{b} \in B(n)} \sum_{i=1}^{n} \frac{b_i}{n} \sum_{j \neq i} f(b_j) \quad = \quad \arg\max_{\boldsymbol{b} \in B(n)} \sum_{i=1}^{n} f(b_i)(n - b_i).$$

Comparing the payoffs of (i)-(iii) yields the desired result. □

E.2 Welfare-maximising equilibria

For any $\Delta \subseteq N$, let $\Delta_{-i} = \Delta \setminus \{i\}$ denote the protection profile where all nodes in Δ different from node i protect. Furthermore, let

$$\begin{aligned} h_i(G, \Delta_{-i}) &= U^i(G, \Delta_{-i} \cup \{i\}) - U^i(g, \Delta_{-i}), & (72) \\ H_i(G, \Delta_{-i}) &= U^{\mathsf{D}}(G, \Delta_{-i} \cup \{i\}) - U^{\mathsf{D}}(g, \Delta_{-i}). & (73) \end{aligned}$$

In words, $h_i(G, \Delta_{-i})$ and $H_i(G, \Delta_{-i})$ are, respectively, the gain of node i and of the designer from i protecting under network G and defence profile Δ_{-i}. Recall that $C_i(G)$ denotes the component of G such that $i \in C_i(G)$. Thus, $C_i(G - \Delta_{-i})$ denotes the set of unprotected nodes in G which have a path to i through unprotected nodes. We can therefore write

$$h_i(G, \Delta_{-i}) = \frac{1}{n}\left[\frac{f(|C_i(G)|)}{|C_i(G)|} - 0\right] + \sum_{j \in C_i(G-\Delta_{-i}) \setminus \{i\}} \frac{1}{n}\left[\frac{f(|C_i(G - E_j(G|\Delta_{-i} \cup \{i\}))|)}{|C_i(G - E_j(G|\Delta_{-i} \cup \{i\}))|} - 0\right] - c. \quad (74)$$

The following lemma establishes that, due to positive externalities, there can never be over-investment in protection.

Lemma 14. *For any G and Δ, $H_i(G, \Delta_{-i}) \geq h_i(G, \Delta_{-i})$, with strict inequality if and only if $|C_i(G)| \geq 2$.*

Proof. Note that

$$\begin{aligned} H_i(G, \Delta_{-i}) &\geq \frac{1}{n}\left[f(|C_i(G)|) - 0\right] + \sum_{j \in C_i(G-\Delta_{-i}) \setminus \{i\}} \frac{1}{n}\left[\frac{f(|C_i(G - E_j(G|\Delta_{-i} \cup \{i\}))|)}{|C_i(G - E_j(G|\Delta_{-i} \cup \{i\}))|} - 0\right] - c, \\ &\geq h_i(G, \Delta_{-i}). \end{aligned}$$

This establishes the first statement in the lemma. For the second statement, consider the direction from right to left. Since f is increasing, the second inequality is strict if $|C_i(G)| \geq 2$. Finally, consider the direction from left to right. If $|C_i(G)| = 1$, then $C_i(G - \Delta_{-i}) \setminus \{i\} = \varnothing$. In this case, $h_i(G, \Delta_{-i}) = \frac{1}{n}\left[\frac{f(1)}{1} - 0\right] = \frac{1}{n}\left[f(1) - 0\right] = H_i(G, \Delta_{-i})$. □

As a corollary, we have that if the first best features no protection, then there is no cost of decentralization.

Corollary 3. *Let (G, Δ) be first best. If $\Delta = \varnothing$, then Δ is an equilibrium of $\Gamma(G)$.*

Proof. For a contradiction, suppose that in the first best the designer chooses G and $\Delta = \varnothing$, but $\Delta = \varnothing$ is not an equilibrium of $\Gamma(G)$. Set $\Delta_{-i} = \varnothing$. It must be that $h_i(G, \Delta_{-i}) > 0$. Since $H_i(G, \Delta_{-i}) \geq h_i(G, \Delta_{-i})$, $\Delta = \varnothing$ cannot be first best. □

The next lemma shows that if $c > t_n(n)$ then D chooses an optimal unprotected network.

Lemma 15. *Let (G, Δ) be an equilibrium of Γ. If $c > t_n(n)$, then $\Delta = \varnothing$ and G is an optimal unprotected network.*

Proof. It suffices to show that, if $c > t_n(n)$, then for any G the unique equilibrium of $\Gamma(G)$ is $\Delta = \varnothing$. To see this, note that, for any $G \in \mathcal{G}(N)$ and $\Delta \subseteq N$, $h_i(G, \Delta_{-i}) \leq \frac{f(n)}{n^2} + (n-1)\frac{1}{n}\frac{f(n-1)}{(n-1)} - c$. Thus, if $c > t_n(n)$, then $h_i(G, \Delta_{-i}) < 0$ for any G and Δ. Not to protect is a dominant strategy for a node on any network, and so for any $G \in \mathcal{G}(N)$, the unique equilibrium of $\Gamma(G)$ is $\Delta = \varnothing$. □

We first prove Proposition 10 (for general f), and then prove Proposition 8.

Proof of Proposition 10. We address Cases 1-3 separately.

Case 1. $t_u(n) < c \leq \min\{t_{u+1}(n), \hat{c}_1(n), \hat{c}_2(n)\}$ for some $u = 0, \ldots, n-1$. Since $c \leq \min\{\hat{c}_1(n), \hat{c}_2(n)\}$, first best is full protection in a connected network. The following claim states that if $c > t_u(n)$ then, for any G, at least u nodes are unprotected in every equilibrium of $\Gamma(G)$.

Claim 1. Suppose $c > t_u(n)$. For any G, $|N \setminus \Delta| \geq u$ in every equilibrium of $\Gamma(G)$.

Proof. Suppose that $c > t_u(n)$ but there exists G such that Δ is an equilibrium of $\Gamma(G)$ and $u' = n - |\Delta| < u$ nodes are unprotected. For any protected node $i \in \Delta$, note that

$$
\begin{aligned}
h_i(G, \Delta_{-i}) &\leq \frac{f(n)}{n^2} + u'\frac{f(n-1)}{n(n-1)} - c < \frac{f(n)}{n^2} + u'\frac{f(n-1)}{n(n-1)} - t_u(n) \\
&= [u' - (u-1)]\frac{f(n-1)}{n(n-1)} \leq 0.
\end{aligned}
$$

Therefore, any node $i \in \Delta$ would rather unprotect, a contradiction. □

Maximum equilibrium welfare is therefore achieved if there are exactly $|N \setminus \Delta| = u$ unprotected nodes such that if an unprotected $i \in N \setminus \Delta$ is attacked, the attack neither spreads nor disconnects G. Equilibrium welfare is therefore bounded above by $\frac{u}{n}f(n-1) + \frac{n-u}{n}f(n) - (n-u)c$. The following claim establishes that G attains maximum equilibrium welfare if and only if $G \in \mathcal{G}^{n-u}(N)$, and thus completes the proof of Case 1.

Claim 2. Suppose $t_u(n) < c \leq t_{u+1}(n)$. There exists an equilibrium Δ of $\Gamma(G)$ such that $U^{\mathsf{D}}(G, \Delta) = \frac{u}{n}f(n-1) + \frac{n-u}{n}f(n) - (n-u)c$ if and only if $G \in \mathcal{G}^{n-u}(N)$.

Proof. For the direction right to left, pick a network $G \in \mathcal{G}^{n-u}(N)$. For a set of nodes $U \subseteq N$ satisfying the conditions for $G \in \mathcal{G}^{n-u}(N)$, consider the defence profile $\Delta = N \setminus U$. For $i \in \Delta$,

$$
h_i(G, \Delta_{-i}) = \frac{f(n)}{n^2} + u\frac{f(n-1)}{n(n-1)} - c \geq \frac{f(n)}{n^2} + u\frac{f(n-1)}{n(n-1)} - t_{u+1}(n) = 0, \quad (75)
$$

and so protected nodes do not wish to deviate. If $u = 0$, then Δ is an equilibrium of $\Gamma(G)$ and the statement is true. Suppose $u \geq 1$. For $j \in N \setminus \Delta$,

$$h_i(G, \Delta_{-i}) \;=\; \frac{f(n)}{n^2} - c < \frac{f(n)}{n^2} - t_u(n) = (u-1)\frac{f(n-1)}{n(n-1)} \leq 0. \tag{76}$$

Combining (75) and (76), we conclude that Δ is an equilibrium of $\Gamma(G)$, and it achieves welfare $U^{\mathbf{D}}(G, \Delta) = \frac{u}{n} f(n-1) + \frac{n-u}{n} f(n) - (n-u)c$.

Consider next the direction left to right. For a contradiction, suppose that there exists $G \notin \mathcal{G}^{n-u}(N)$ which achieves maximum equilibrium welfare. If $u = 0$, then $G \notin \mathcal{G}^{n-u}(N)$ means that G is not connected. Since $c \leq t_1(n)$, full protection is the unique equilibrium on every network. Thus, the designer can be strictly better off by choosing a connected network, a contradiction.

Suppose then that $u \geq 1$. Since G attains maximum equilibrium welfare, there are only u unprotected nodes whose potential attack neither spreads nor disconnects the network. Since the attack to $i \in N \setminus \Delta$ does not to spread, it must be that $ij \notin G$ for all $j \in N \setminus \Delta$. Since the attack to $i \in N \setminus \Delta$ does not disconnect the network, it must be that $G - \{i\}$ is connected. It follows from these two observations that $G \notin \mathcal{G}^{n-u}(V)$ implies that there exists a pair of nodes (i, j), $i \in N \setminus \Delta$ and $j \in \Delta$, such that $ij \notin G$. Then, for node j,

$$h_j(G, \Delta_{-j}) \;\leq\; \frac{f(n)}{n^2} + (u-1)\frac{f(n-1)}{n(n-1)} - c < \frac{f(n)}{n^2} + (u-1)\frac{f(n-1)}{n(n-1)} - t_u(n) = 0.$$

That is, node j would strictly prefer not to protect, a contradiction. □

Case 2. $\hat{c}_1(n) < c \leq \min\{\hat{c}_3(n), t_n(n)\}$. Since $\hat{c}_1(n) < c \leq \hat{c}_3(n)$, first best payoffs of the designer are attained by a centre-protected star. It is easy to check that, since $c \leq t_n(n)$, the centre m of the star protects if no spoke protects. If a spoke finds protection profitable when m protects, then by Lemma 14 the designer would be strictly better off, and therefore full protection would be optimal, a contradiction. Hence $\Delta = \{m\}$ is an equilibrium of the star and the designer achieves first best payoffs.

Case 3. $c \leq \min\{t_n(n), \max\{\hat{c}_2(n), \hat{c}_3(n)\}\}$. If $c > \max\{\hat{c}_2(n), \hat{c}_3(n)\}$, then first best is an optimal unprotected network. By Corollary 3 this is attainable in equilibrium. If $c > t_n(n)$, then by Lemma 15 $\mathbf{D}$ chooses an optimal unprotected network. □

Before proving Proposition 8, we show that if $f(y) = y^2$ then the optimal unprotected network consists of two components, of sizes $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$.

Lemma 16. *Assume $f(y) = y^2$ and suppose the attack is random. If (G, Δ) is first best and $\Delta = \varnothing$, then G consists of two components, of sizes $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$.*

Proof. Let $\boldsymbol{b}^*$ be an optimal partition, i.e.

$$\boldsymbol{b}^* \in \arg\max_{\boldsymbol{b} \in B(n)} \sum_{i=1}^{n} h(b_i),$$

where $h(b_i) = (n - b_i)b_i^2$. Note that $h'(b_i) = 0 \Leftrightarrow b_i = \frac{2n}{3}$, and $h''(b_i) = 0 \Leftrightarrow b_i = \frac{n}{3}$. That is, function $h(\cdot)$ has a maximum at $\frac{2n}{3}$, and is convex on $\left[0, \frac{n}{3}\right]$ and concave on $\left[\frac{n}{3}, n\right]$. We show that the optimal partition contains two components, of sizes $\lceil n/2 \rceil$ $\lfloor n/2 \rfloor$, with the following steps.

Step 1. $b_1 < \left\lceil \frac{2n}{3} \right\rceil + 1$. For a contradiction, suppose that $b_1 \geq \left\lceil \frac{2n}{3} \right\rceil + 1$. Then consider the partition which is equal to $\boldsymbol{b}$ except that we isolate one node from b_1. Since $h(\cdot)$ is decreasing on $[\frac{2n}{3}, n]$ and increasing otherwise, this is a strict improvement, a contradiction.

Step 2. There is at most one component of size b_i such that $0 < b_i < \left\lfloor \frac{n}{3} \right\rfloor$. Suppose $b_i \leq b_j < \left\lfloor \frac{n}{3} \right\rfloor$. Then consider moving nodes from the subset of size b_i to the subset of size b_j, up to the point in which the new size of the larger subset is $b'_j = \min\left\{\left\lfloor \frac{n}{3} \right\rfloor, b_j + b_i\right\}$. Since $h(\cdot)$ is convex on $\left[0, \frac{n}{3}\right]$, this is a strict improvement, a contradiction.

Step 3. If $b_i, b_j \geq \left\lceil \frac{n}{3} \right\rceil$, then $|b_i - b_j| \leq 1$. If $|b_i - b_j| > 1$, then move one element from the larger to the smaller subset. Since $h(\cdot)$ is concave on $\left[\frac{n}{3}, n\right]$, this is a strict improvement, a contradiction.
By steps 1-3, there are two possibilities. The first possibility is that there are only two non-empty subsets, of sizes $b_i > \left\lfloor \frac{n}{3} \right\rfloor$ and $b_j \leq \left\lfloor \frac{n}{3} \right\rfloor$. It is easy to verify that the optimal partition into two components is with sizes $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$, a contradiction. Since $a = n/3$, the second possibility is that there are two subsets of sizes greater than or equal to $\left\lfloor \frac{n}{3} \right\rfloor$, and possibly one subset of size less than or equal to $\left\lfloor \frac{n}{3} \right\rfloor$. Let x, y, z denote the sizes of the three components, with $x \geq y \geq 1/3 \geq z \geq 0$ and $x + y + z = n$. Abstracting from integer problems, maximising $(n - x)x^2 + (n - y)y^2 + (n - z)z^2$ with respect to these constraints yields two constrained local optima: $(x, y, z) = (n/3, n/3, n/3)$ and $(x, y, z) = (n/2, n/2, 0)$. It is straightforward to verify that the objective is maximised in the latter. Hence, the optimal unprotected network has two components, of sizes $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$. □

Proof of Proposition 8. It follows from Proposition 10 and Lemma 16. □

E.3 Welfare-minimizing equilibria

Proof of Fact 1. Let G be a connected network. Pick any node $i \in N$. For any $\Delta \subseteq N$, note that $h_i(G, \Delta_{-i}) \geq \frac{f(n)}{n^2} - c \geq 0$, where the last equality uses $c \leq t_1(n) = \frac{f(n)}{n^2}$. Therefore, the unique equilibrium of $\Gamma(G)$ is $\Delta = N$. □

Proof of Fact 2. Consider statement (1) first. Suppose G is not complete and attains maximum equilibrium welfare in every equilibrium of $\Gamma(G)$. By Proposition 10, every equilibrium Δ of $\Gamma(G)$ must be such that there is exactly one unprotected node l, where $G - \{l\}$ is connected. Fix such an equilibrium Δ. Note that $li \in G$ for every $i \neq$

l. Otherwise there would be a node i whose neighbors are all protected in Δ, and so $h_i(G, \Delta_{-i}) = \frac{f(n)}{n^2} - c < 0$, a contradiction. Moreover, since G is not the complete network, there exist nodes u_1, u_2 such that $u_1, u_2 \notin G$. To complete the proof, we will show that there is an equilibrium of $\Gamma(G)$ where nodes u_1 and u_2 do not protect, and therefore G does not attain maximum equilibrium welfare in every equilibrium of $\Gamma(G)$. We will use the following concepts.

Definition 2. A set of nodes $C \subseteq N$ is a *vertex cover* (VC) of G if, for all $ij \in G$, $ij \cap C \neq \varnothing$. A vertex cover C is *minimal* (MVC) if, for all $D \subsetneq C$, D is not a vertex cover of G.

Consider the following two steps.

Step 1. If C is an MVC of G and $l \in C$, then $\Delta' = C$ is an equilibrium of $\Gamma(G)$. Let C be an MVC of G, and consider defence profile $\Delta' = C$. For $i \notin \Delta'$, note that $j \in \Delta'$ for every $j \in \partial_G(i)$, or otherwise Δ' would not be a VC (it would not be covering all links in G). Therefore, $h_i(G, \Delta') = \frac{f(n)}{n^2} - c < 0$. For $j \in \Delta'$, there must exist $i \in \partial_G(j)$ such that $i \notin \Delta'$, or otherwise the VC Δ' is not minimal. Furthermore, since $l \in \Delta'$ and $kl \in G$ for every k, an attack on i does not disconnect G. Therefore, $h_j(G, \Delta') \geq \frac{f(n)}{n^2} + \frac{f(n-1)}{n(n-1)} - c \geq 0$, where the inequality uses $c \leq t_2(n) = \frac{f(n)}{n^2} + \frac{f(n-1)}{n(n-1)}$.

Step 2. There exists an MVC Δ' such that $l \in \Delta'$ and $u_1, u_2 \notin \Delta'$. Construct Δ' as follows. Start with $\Delta'_0 = N \setminus \{u_1, u_2\}$. Since $\{u_1, u_2\} \notin G$, Δ'_0 is a VC. If the VC Δ'_0 is not minimal, remove nodes from Δ'_0 until obtaining an MVC. Node l will be in any such MVC, or otherwise the link $lu_1 \in G$ would not be covered.

Combining steps 1 and 2 completes the proof of statement (1). Consider next statement (2). Let G be the complete network G^c. For a contradiction, suppose there exists an equilibrium Δ of $\Gamma(G^c)$ where $|\Delta| \leq n - 2$. Let $e = |N \setminus \Delta| \geq 2$ denote the number of unprotected nodes. For unprotected node $i \notin \Delta$, $h_i(G^c, \Delta) = \frac{f(n)}{n^2} + \frac{e-1}{n}\frac{f(n-e-1)}{n-e-1} - c \geq \frac{e-1}{n}\frac{f(n-e-1)}{n-e-1} - \frac{f(n-1)}{n(n-1)}$, where the inequality uses $c \leq t_2(n)$. By the condition given in the fact, it is straightforward to see that $h_i(G^c, \Delta) \geq 0$. That is, an unprotected node would prefer to protect, a contradiction. □

Proof of Proposition 9. By Proposition 7 and Lemma 16, we have that

$$\begin{aligned}
\hat{c}_1(n) &= \frac{2n-1}{n}, \\
\hat{c}_2(n) &= \frac{n^2 - (\lfloor n/2 \rfloor^2 + n \bmod 2)}{n}, \\
\hat{c}_3(n) &= (n-1)^2 - \left(\lfloor n/2 \rfloor^2 + n \bmod 2\right),
\end{aligned}$$

so that $\hat{c}_1(n) < \hat{c}_2(n) < \hat{c}_3(n)$. By Proposition 7, first best is full protection if $c \leq \hat{c}_1(n)$, a centre-protected star if $\hat{c}_1(n) < c \leq \hat{c}_3(n)$, and the optimal unprotected network if $c > \hat{c}_3(n)$. Furthermore, $t_1(n) = 1$, $t_2(n) = \frac{2n-1}{n}$, and $t_n(n) = (n-1) + \frac{1}{n}$, so that $0 <$

$t_1(n) < \hat{c}_1(n) = t_2(n) < t_n(n) < \hat{c}_3(n)$. Let (G, Δ) be a welfare minimizing equilibrium. We consider the different cases.

Case (1). $0 < c \leq t_1(n)$ Since $t_1(n) < \hat{c}_1(n)$, by Fact 1 any connected network G attains first best welfare in unique equilibrium.

Case (2). $t_1(n) < c \leq \hat{c}_1(n) = t_2(n)$ By Fact 2, G must be the complete network.

Case (3). $\hat{c}_1(n) = t_2(n) < c \leq t_n(n)$ By Proposition 10, the star network has an equilibrium where only the centre protects. To see that this equilibrium is unique, note first that $c > \hat{c}_1(n)$ implies that any other equilibrium must have the centre unprotected. Let G be the star network, and Δ be a defence profile where $s \in \{0, \ldots, n-2\}$ spokes protect. For spoke $j \notin \Delta$, $h_j(G, \Delta) = \frac{f(n)}{n^2} + \frac{n-s-1}{n} f(1) - c < -\frac{s}{n} \leq 0$, where the inequality uses $c > t_2(n)$. Hence the unique equilibrium of the star is $\Delta = \{m\}$, where m is the centre. **D** chooses the star and attains first best payoffs.

Case (4). $c > t_n(n)$ By Lemma 15, G is the optimal unprotected network.

□