Robust Networks

Sanjeev Goyal * Adrien Vigier†

May 14, 2009
First draft: July 2008

Abstract

A network is said to be robust if it performs well against attacks.

We first study random uniform attacks. Robust networks consist of equal size groups whose number increases in the intensity of the attack. We then turn to intelligent attack: the adversary observes the network and chooses nodes to attack. Optimal attack involves targeting only a few nodes and ignoring the rest. In response, robust networks consist of equal size groups. Their number is higher and performance poorer as compared to the case of random attack.

We extend the framework to allow for defence of nodes. It is attractive to protect central nodes since they minimize the prospects of indirect detection/infection. With limited defence budget it is best to minimize the number of such central nodes: the robust network is a star.

*Faculty of Economics & Christ's College, University of Cambridge.

†Faculty of Economics & Queens College, University of Cambridge.

We are grateful to Murali Agastya, Heski Bar-Isaac, Yann Bramoulle, Indranil Chakravarty, Andrea Galeotti, Aditya Goenka, Matthew Jackson, Rufus Pollock, Rakesh Vohra, and seminar participants at Cambridge, Northwestern, NUS and Oxford for comments. Sanjeev Goyal would like to thank Parimal Bag for related discussions on crime and its detection.

1 Introduction

Connections between individuals facilitate the exchange of goods and information but they also expose an individual to threats and dangers faced by the other individuals. In some cases, these threats are of a physical nature, as in the spread of diseases: interacting with other people increases the danger of contacting a disease from them. In other cases, the danger may be virtual: computers may receive viruses and worms from other computers. Finally, individuals dealing in terrorism and crime face the threat of detection and elimination from the police. We say that a network is *robust* if it performs well against attacks.

We start with the following elementary model: a group of individuals/nodes aim to perform a given task. Ex ante the group does not know the task to be performed, and each node is expert in an equal number of tasks. Eventually, the relevant task is revealed to one node in the group. Call this node the informer. If the informer can carry out the task, she does so; if she cannot carry it out then she seeks the relevant expert in the group. Communication operates through a network of interaction. A node may communicate with those she has a link to, but it can also use neighboring nodes as intermediaries to contact their neighbors and so on.

The nodes in the network face attacks: if a node is attacked and 'infected/detected', it is ineffective both for communication as well as for tasks. Moreover, detection renders neighboring nodes and the neighbors of neighbors vulnerable as well.

We shall focus on the situation where, in the absence of attack, it is best to connect all the nodes. In the presence of an adversary, indirect detection of nodes implies a trade-off in the decision to connect nodes and we explore how this affects robust networks.

We start by characterizing optimal networks under uniform random detection of every node. We show that optimal networks consist of equal size components and that the component size is declining in the probability of detection (Proposition 1). For example, in a setting with 4 nodes, the optimal network contains a single component when this probability is small, then two equal size components, followed by four isolated individuals, when the probability is large.

The key aspect of this result is that unequal size groups cannot be optimal. This is an interesting and difficult property to establish. The natural way to prove this property would be to show that in a class of networks with k groups, k equal groups dominate networks with unequal groups. However, this is NOT true in general: Figure 1 presents an example in which, for a range of attack levels, a network with two groups containing 3 and 1 nodes, respectively,

dominates a network with two groups of equal size. Our argument has to show that when a network with k unequal groups dominates a network with k equal groups then there is a network with $k' \neq k$ equal groups which dominates the network with k unequal groups.

We next study robustness when an intelligent adversary observes the network and can target specific nodes. We show that the optimal attack strategy of the adversary is *asymmetric* and involves targeting a few nodes and ignoring the rest. From the point of view of the designer, robust networks consist of equal size components whose number grows (and size falls) as the attack budget of the adversary increases (Proposition 2). However, there is a sharp contrast in the size and number of groups in the robust network. For instance, with unit budget of attack, as number of nodes grows, in the uniform random attack world the robust network constitutes a single group. By contrast, the intelligent adversary always targets a single node and the optimal network consists of two equal components, irrespective of the number of nodes! These results show how incorporating an intelligent adversary leads to very different predictions both with regard to attack strategies, robust network architecture as well as the performance attainable.

In actual practice, individuals, firms and countries, who have an interest in the smooth functioning of the network invest resources to protect the network. So the robustness of the network is defined out of interaction between these defenders and adversaries of the network. This motivates an extension of the model where we enrich the strategic options of the designer: he designs the network *and* protects nodes in the network. Defence of a node helps protect the functionality of the node – to communicate and do tasks – but it also serves another important new function: it can limit indirect detection of nodes. Thus it is attractive to protect central nodes. If the defence budget is small, this in turn makes it attractive to reduce the number of central nodes: a star network with a protected center is robust (Propositions 3-5).

Empirical work has highlighted the salience of highly connected hub nodes in real world networks; for a survey of this work see Goyal (2007) and Barabasi (1999). Many of these networks – such as internet and transport networks – face intelligent adversaries. In an influential paper, Albert, Jeong, Barabási (2000) argue that networks containing hubs – the so called scale-free networks – are robust to random attacks but very vulnerable to targeted attacks. Our results stand in sharp contrast and highlight the value of a strategic approach founded upon individual incentives: in a world where the designer can choose a network and protect specific nodes while an intelligent adversary can target specific nodes, networks with hubs are robust.

There is a large literature on communication structures in economics; see e.g., Bolton and

Dewatripont (1994), Radner (1992, 1993) and van Zandt (1999), and Garicano (2000). However, there have been very few attempts at developing a strategic approach to communication network design in the face of an intelligent adversary. To the best of our knowledge, the only exception is Baccara and Bar-Issac (2007). They consider a sequential move model in which the legal authority commits to a monitoring strategy while the organization then chooses an optimal communication design. The members of the organization are engaged in a infinitely repeated game and the links between individuals transmit personal information about agents which facilitates the implementation of punishments in the repeated game. In their model, links are directed and represent 'power' over agents. By contrast, in our model, links are undirected and serve to communicate information about tasks. Moreover, we consider a model with design as well as defence of nodes. So the methods of analysis and the results are quite different.

In a recent paper, Hong (2008) investigates the strategic complementarities between linking and protection in a game of network formation and protection played among nodes. Similarly, Bala and Goyal (2000) study a game of network formation among nodes faced with an exogenously given uniform probability of link deletion. By contrast, in the present paper, the focus is on strategic interaction between a network designer and an intelligent adversary.

There is also a very large literature on network security spread across disciplines such as computer science, statistical physics, engineering and operations research (Barabasi (1999); Nagaraja and Anderson (2007); Smith (2008); Levine (1999)). Most of this work studies network security using a single agent or planner's optimization perspective. In some of the models the interest is in the design of survivable networks, see e.g., Grotschel, Monma and Stoer (1995), while in others the focus is on the adversary's optimal network attack strategy, see e.g., Smith (2008). By contrast, Nagaraja and Anderson (2007) take an explicitly strategic approach to network security. They present simulation results on the evolution of defence and attack strategies over time. The literatures are vast, but to the best of our knowledge, the equilibrium analysis of a game in which a designer chooses the network (and protection) while an adversary chooses nodes to target is novel.

The rest of the paper is organized as follows. Section 2 presents the basic model. In section 3 we study the pure network design problem, while in section 4 we allow the designer to choose the network and also defend some nodes in the network. Section 5 discusses some extensions of the model to allow for general payoffs and imperfect indirect detection. Section 6 concludes. All the proofs are presented in the Appendix.

2 A simple model of communication

Let $N = \{1, ..., n\}$ denote a set of n nodes, and let $N \times N$ be the set of states of the world. Let $(\Omega, \mathcal{F}, P)$ denote a probability triple, with $X : \Omega \to N$, and $Y : \Omega \to N$ two independent uniform random variables. On the event $\{\omega \in \Omega; X(\omega) = i, Y(\omega) = j\}$ we say that node i is the expert, while node j is the informer.

Nodes in N communicate through a network of interaction g. A network g for the set of nodes N is represented by an $n \times n$ matrix, where $g_{ij} = 1$ if there is a link between i and j and $g_{ij} = 0$ otherwise. Let $\mathcal{G}$ denote the set of graphs defined on N. We assume that $g_{ii} = 1$, $\forall i \in N$, $\forall g \in \mathcal{G}$, so that every node is connected to itself. Links are also assumed to be undirected, i.e. $g_{ij} = g_{ji}$, $\forall i, j \in N$, $\forall g \in \mathcal{G}$. We say that there is a path between two nodes i and j in the graph g if there exists a sequence of nodes $i_1,.., i_k$ such that $g_{ii_1} = g_{i_1 i_2} = ... = g_{i_{k-1} i_k} = g_{i_k j} = 1$. Two nodes are connected if and only if there is a path between them. A component of the graph g is a maximal connected subset. The set of components of the network g is denoted by $\mathcal{C}(g)$. Notice that $\mathcal{C}(g)$ provides a partition of the set N. We assume that any connected pair of nodes is capable of communication.

We now discuss the potential obstacles to communication resulting from attack by an adversary. Let $(Q_i)_{i \in N}$ denote a set of Bernoulli random variables $\Omega \to \{0, 1\}$ independent of each other and of the state of the world. Let $\mathbf{q} \in [0, 1]^n$ denote the mean of the random variables $(Q_i)_{i \in N}$, i.e. $\mathbf{q} = (P(Q_1 = 1), .., P(Q_n = 1))$. We shall call q_i the attack/monitoring intensity of i. We assume that on the event $\{\omega \in \Omega; Q_i(\omega) = 1\}$ all nodes connected to i are prevented from communication; thus indirect detection and elimination is perfect. This assumption is a mirror image of the assumption that communication across paths is perfect.

These assumptions on perfect communication and indirect detection are strong; but they appear to be a natural point to start the analysis of robust networks. In section 5 we briefly discuss more general communication and indirect detection in networks.

Let $V(g, \mathbf{q})$ denote the probability that informer and expert can communicate in network g under monitoring profile $\mathbf{q}$. We shall call $V(g, \mathbf{q})$ the probability of task completion for the group N. Let $f(m)$ denote the probability of task completion for a component of size m conditional upon $Q_i = 0$ for all i in that component. Thus,

$$f(m) = \left(\frac{m}{n}\right)^2 \tag{1}$$

Notice that $f(m)$ is increasing and convex in m; these properties play a key role in our analysis.

Communication within $C \in C(g)$ is possible if and only if $Q_i = 0$, for all $i \in C$. This event occurs with probability $\prod_{i \in C}(1 - q_i)$. Summing across components, we obtain the following simple expression for the probability of task completion:

$$V(g, \mathbf{q}) = \sum_{C \in C(g)} \left[\frac{\#C}{n}\right]^2 \prod_{i \in C}(1 - q_i) \tag{2}$$

In particular, $V(g, \mathbf{q})$ is component additive. Thus, in the basic model, $V(g, \mathbf{q})$ hinges solely on the number and size of the components in network g. Section 4 extends the model to allow for the defence of the network by the designer: payoffs in that model do depend on the details of the network structure.

A network made up of two components with size n_1 and n_2 is denoted (n_1, n_2). A network made up of k components of equal size is called k-balanced and denoted by g^{bk}, i.e. $g^{bk} = (\frac{n}{k}, .., \frac{n}{k})$. In particular, the empty network is denoted by $(1, .., 1)$ while (n) is used to represent the complete network. We let $\mathcal{G}_k \subset \mathcal{G}$ denote the subset of graphs with at most k components.

We shall say that a network g is *robust* if it maximizes the probability of task completion for the group N.

3 Random versus strategic attack

We start with an analysis of the network design problem for symmetric detection probability, $q_i = q$, for all $i \in N$. This is the natural benchmark as it allows us to understand the implications of strategic attack. This model is also of some independent interest as it offers an analysis of robust networks in contexts of attack which are biological or physical.

To get an idea of the different factors at work here, consider $q = 0$, and $q = 1$. If $q = 0$, then a single component is optimal by convexity of f. While if q is close to 1 on the other hand, inspection of (2) shows that $V(g, \mathbf{q})$ may be reduced to the contribution of its smallest components. The empty network is therefore optimal under high monitoring. In general, when $q > 0$, separating out nodes may in fact enhance the prospects for communication among nodes due to the effect of indirect detection. Consider, by way of illustration, the reduced problem with equal size components and let k denote the number of components in network g. It is easily checked that

$$V(g^{bk}, q) = \frac{1}{k}(1 - q)^{\frac{n}{k}}$$

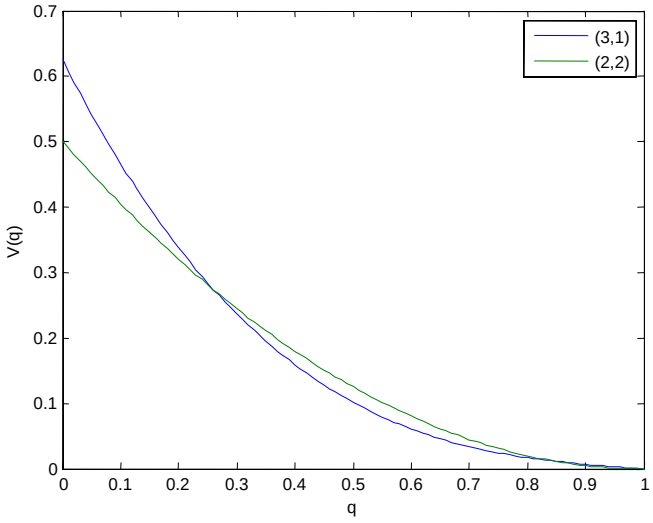

Figure 1: Equal versus unequal groups

and that the unique solution to the optimization problem, k^*, satisfies $\frac{dk^*}{dq} > 0$. The optimal component size therefore diminishes as q increases.

In general, however, unbalanced networks may be optimal and we therefore need to consider this possibility. Figure 1 compares probabilities for two networks $(2, 2)$ and $(3, 1)$ and illustrates some of the difficulties involved in addressing this problem. For small values of q the effect of convexity dominates, and the unbalanced network (3,1) is best due to the influence of the large component. For large values of q on the other hand monitoring has the stronger effect, and the unbalanced network (3,1) is best due to the influence of the small component. For intermediate values of q the balanced network is efficient. Thus, it is not true that a network with two equal groups dominates networks with two unequal groups, for all values of attack q.

What can we say for networks in general? Figure 2 shows that, for small q, (4) dominates $(3, 1)$ whenever the latter dominates $(2, 2)$. Similarly, at q large, $(1, 1, 1, 1)$ dominates $(3, 1)$ whenever the latter dominates $(2, 2)$. Hence, the unbalanced network is never optimal.

The following result shows that these considerations hold generally.

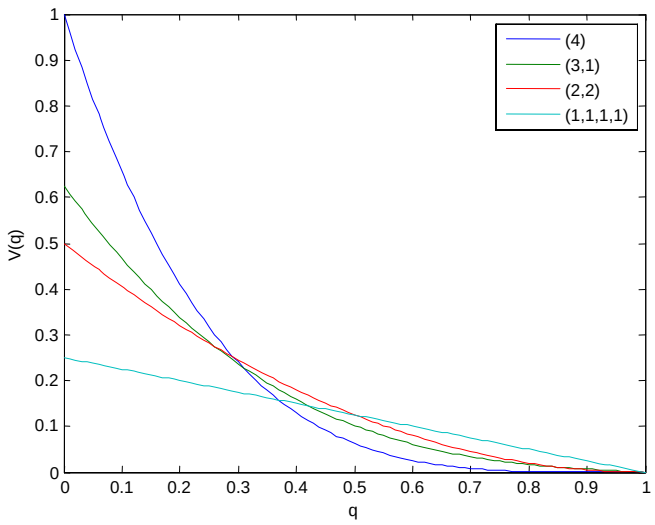

Figure 2: Robust networks: $n = 4$, $q \in [0, 1]$

Proposition 1 *Suppose payoffs are given by (2), the number of nodes n is even and there is uniform random attack* $\mathbf{q} = (q, .., q)$. *There exist* $\widetilde{q}$, η , $\zeta > 0$ *such that: (i)* (n) *is uniquely robust for* $q \in [0, \widetilde{q})$, *where* $\widetilde{q} \leq 1/2$. *(ii)* $(\frac{n}{2}, \frac{n}{2})$ *is uniquely robust for* $q \in (\widetilde{q}, \widetilde{q} + \eta)$. *(iii)* $(2, .., 2)$ *is uniquely robust for* $q \in (\frac{1}{2} - \zeta, \frac{1}{2})$. *(iv)* $(1, .., 1)$ *is uniquely robust for* $q \in (\frac{1}{2}, 1]$. *(v)* $\forall q \in [0, 1]$, $\forall g(q)$ *robust given q, no group in* $C(g(q))$ *has t times the size of another,* $t \in \{2, 3, ..\}$.

Figure 3 illustrates payoffs from the different networks for $n = 4$, across all $q \in [0, 1]$. A connected network is robust for $q \in [0, 0.28]$, a network with two equal components is robust for $q \in [0.28, 0.5]$, while a network with four isolated nodes is optimal for $q \in [0.5, 1]$.

We now turn to the case of strategic attack: the adversary can target specific nodes based on its knowledge of the network and the designer chooses a network in anticipation of this intelligent attack. We define a game of perfect information between a network designer $\mathcal{D}$ and her adversary $\mathcal{A}$. Let $a \in \mathbb{N}$ denote the budget of the adversary. $\mathcal{D}$ moves first and chooses $g \in \mathcal{G}$. The adversary $\mathcal{A}$ observes g and chooses a monitoring profile $\mathbf{q} \in [0, 1]^n$, such that $\sum_i q_i = a$. Notice that our formulation equates resources and probability of successful attack; this formulation is analytically convenient and is borrowed from Baccara and Bar-Isaac (2008).

To see how this formulation may be motivated suppose that the adversary allocates resources r_i to specific nodes and that these resources translate to a probability of successful deletion of the specific node via a function $e(r_i)$ which takes on values between 0 and 1. Our present model then refers to the case where $r_i \in [0,1]$ and $e(r_i) = r_i$.[1]

Payoffs of the two players are given by

$$\pi^{\mathcal{D}} = -\pi^{\mathcal{A}} = V(g, \mathbf{q}) \tag{3}$$

We refer to this game as the DA game; in this section, we shall say that a network $g \in \mathcal{G}$ is *robust* if there exists a sub-game perfect equilibrium of the DA game which supports it.

The following result characterizes robust networks in the DA game.

Proposition 2 *In the DA game: if* $n \bmod 2a = 0$ *the unique robust network consists of* $2a$ *equal size components; if* $n \bmod 2a \neq 0$*, then all components, except possibly one, have equal size; if* $a > \frac{n}{2}$*, the empty network is uniquely robust.*

Figure 3 illustrates this result for $n = 12$ and $a = 1, 2, 3, 6$.

First, observe that, in equilibrium, at most one node is monitored in any component. When $\mathcal{A}$ monitors two nodes part of her effort is redundant since, with strictly positive probability, she detects both nodes independently with no additional payoffs. Second, in equilibrium, $\mathcal{A}$ will target only the largest a components. But this means, in turn, that the a largest components must have equal size. If one of them is strictly larger, then the designer $\mathcal{D}$ can strictly increase her payoffs by removing one node out of the largest component. Finally, observe from the convexity of f, that non-targeted components must also have maximal size. We have thus proved that, modulo integer constraints, robust networks consist of equal size groups. Let this group size be $s = n/k$, where k is the number of groups. The payoff to the designer is

$$\left(\frac{s}{n}\right)^2 \left(\frac{n}{s} - a\right) \tag{4}$$

This payoff is quadratic in s with a unique maximum attained at $s = \frac{n}{2a}$.

Proposition 2 yields a number of insights. One, optimal attack strategies for $\mathcal{A}$ involve targeting a and ignoring the other $(n-a)$ nodes. For $n >> a$ the attack strategy of an intelligent adversary who observes the network is thus strikingly asymmetric and in sharp contrast

[1] Our main finding on optimal attack strategy is that it targets a few nodes and ignores the rest (see Proposition 2 below). This result is reinforced if $e(.)$ is convex. On the other hand, if $e(.)$ is concave, there are diminishing returns and targeting a few nodes (and ignoring the rest) may then no longer be optimal.

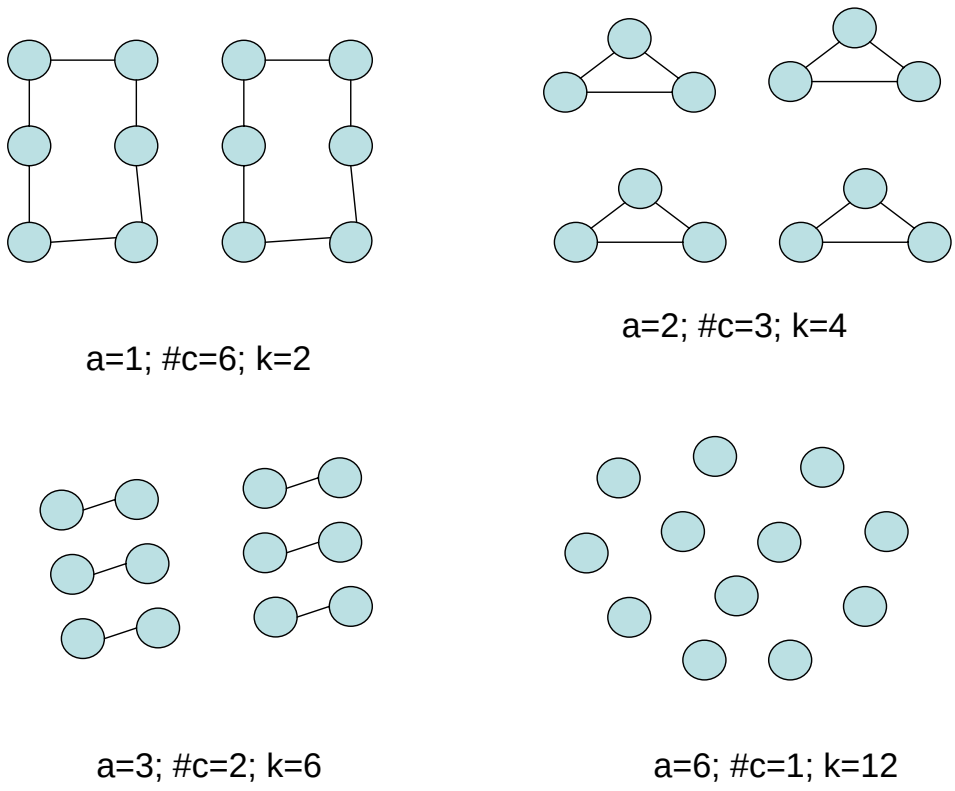

Figure 3: Robust networks: n=12, varying attack budgets

to the uniform random attack studied in the previous section. Two, in spite of the different pattern of attack, optimal networks remain 'balanced', with the optimal cell size declining in the adversary's budget. At a superficial level, this is similar to the architecture of optimal networks in response to uniform random attack. However, the quantitative implications of the results are very different. The following example illustrates this difference.

Example 1 *Random vs targeted attack.*

Fix $a = 1$ and raise n. Proposition 1 says that under random attack, eventually the optimal network contains one component and is connected. The probability of task completion is $(1 - \frac{1}{n})^n \sim e^{-1} \sim \frac{1}{2.72}$. Proposition 2, on the other hand, says that under strategic attack robust networks have two components and the adversary targets at most one node in each of them. The probability of task completion is 1/4, *irrespective* of the number of nodes. Thus the number of components in the robust network differ (1 versus 2) and the probability of task completion is also very different (0.4 versus 0.25) when we compare random attack with strategic attack. $\triangle$

4 Network Design and Defence

In many contexts of practical interest, the designer can protect a subset of nodes in addition to designing the network. What is a robust network in such a situation? To address this question, we study a game in which $\mathcal{D}$ first chooses a network. $\mathcal{D}$ and $\mathcal{A}$ then play a game of attack and defence on the chosen network.

The key issue in such a game is to describe how the attack and defence of a node interact to determine its status and how does defence of a node alter the likelihood of its indirect detection.

Let us first take up the status of a node which is attacked by the adversary and also defended by the designer. A natural assumption is that attack and defence are equally effective: so if 1 unit of defence and attack resources are targeted on one node then there is one half probability that it is successfully defended and one half probability that it is eliminated.[2] Similarly, if the adversary (designer) targets a node with 1 unit of resource and the designer (adversary) chooses not to defend (not to attack) this node then the node is eliminated (successfully defended) with probability 1.

Consider next the issue of indirect detection, where node i is defended by the designer and some other node j in the same component is attacked by the adversary. For simplicity, we shall assume that a defended node is immune to indirect detection/elimination.[3] These considerations are summarized in the following assumption.

Assumption A.1: *Attack and defence decisions are indivisible. If a defended node is attacked, then it is eliminated with probability 1/2. Elimination of a node leads to infection and elimination of nodes on paths leading away from this node until the process encounters a successfully defended node. A defended node is immune to indirect detection.*

Thus defended nodes act as fire-walls and prevent the spread of infections/eliminations. This role of defence is a key element in our analysis. Figure 4 illustrates some aspects of defence.

We consider all three possibilities with regard to the order of moves of the designer and the adversary. In the designer first version, $\mathcal{D}$ first chooses a node to defend, which is observed by $\mathcal{A}$ who then chooses a node to attack. The adversary first version is similar, but with

[2]The standard model of contests due to Tullock (1967), offers a natural foundation for this assumption.

[3]A weaker assumption would be that the probability of indirect detection is positive but small. It is possible to extend our arguments in Proposition 3 and 4 to allow for small but positive probability of indirect detection.

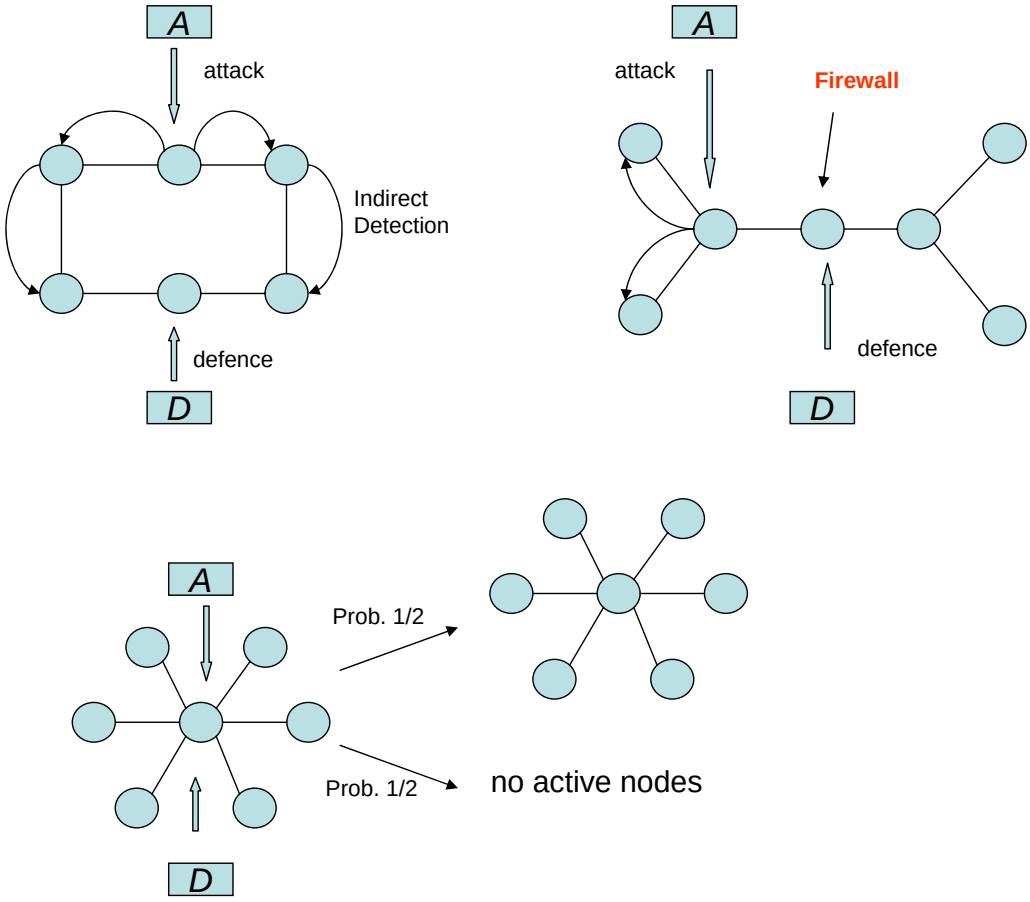

Figure 4: Attack, defence and fire-walls

the order of moves reversed. Finally, in the simultaneous game, players choose which node to attack and defend without observing the other's action (or at the same time).

As before, a denotes the attack budget of the adversary. We let d denote the defence budget of the designer. We start with the case in which attack and defence budgets are equal to 1. Recall that in the basic model ($d = 0$), a robust network consists of two equal size components. We therefore restrict attention to networks in $\mathcal{G}_2$ in the first instance. The following result characterizes robust networks with attack and defence for *all* order of moves in the attack and defence game.

Proposition 3 *Consider the DA game with defence. Suppose* ***A.1*** *holds and let* $g \in \mathcal{G}_2$*. If* $a = 1 = d$ *then, for every order of moves, the star is robust.*

This result builds on a general property of graphs: there exists a node (say) i with the property that for any node $j \neq i$, there exist j-independent paths between i and at least half the nodes in the graph. This property of graphs has the following important implication. Independently of the order of moves the adversary can guarantee herself the payoff $-\frac{1}{2}f(n)$. However, this is exactly the payoff that the adversary earns if the designer chooses a star network and defends the central node. This means that a single star is robust within the set of single component networks. The argument is completed by showing that the maximum payoffs to the designer in networks with two components is no more than $\frac{1}{2}f(n)$.

A comparison of Propositions 2 and 3 illustrates in a dramatic manner the effects of node defence on the architecture of robust networks. In the pure design setting, robust networks consist of two components and the architecture of each component is not determined. By contrast, if the designer can defend a node then the robust network is a single star and the designer defends the central node. Thus defence alters the number of components in the network as well as the architecture within the component. We now explore the scope of this result by allowing larger budgets of attack for the adversary. For simplicity, we restrict attention to the designer move first game. We will retain the assumption that the defended node is immune to indirect detection. However, the possibility of larger budgets raises an issue about whether a node can be attacked by more than 1 unit of resource. In the basic model, this was not an issue as 1 unit of resource sufficed to ensure its elimination. However, in a model with defence, the status of a node depends on the resources allocated for defence *and* for attack. We suppose that the probability of successful defence is given by the familiar Tullock contest function: if the designer allocates d_i units to node i and the adversary allocates

a_i units to the same node then the probability of successful defence is given by $\frac{d_i}{a_i+d_i}$. The assumptions regarding indirect detection remain as in Assumption A.1. We state:

Assumption A.1′: *Attack and defence decisions are indivisible. If a defended node i is attacked with resource a_i, then it is successfully defended (eliminated) with probability $\frac{1}{a_i+1}$ ($\frac{a_i}{a_i+1}$). Successful elimination of a node leads to infection and elimination of nodes on paths leading away from this node until the process encounters a successfully defended node. A defended node is immune to indirect detection.*

Consider connected networks; our next result highlights the important benefits of centralizing communication: a star with protected center is the ideal way to control indirect detection when the designer has limited defence budget.

Proposition 4 *Consider the DA game with defence. Suppose **A.1′** holds and let $g \in \mathcal{G}_1$. If $d = 1$ then, for arbitrary a, the star is robust.*

We have been unable to characterize robust networks in the set of all networks and general attack budgets. However, simple computations show that the payoff from a network with $2a$ equal size groups, is $1/(4a)$ which is clearly smaller than $1/(a+1)$, the payoff to the designer from the star with protected center. We complete this Section by noting that the proof of Proposition 4 is easily adapted to show that, for $d = 1$ and under Assumption **A.1′**, the star is also robust under the kind of uniform random attack studied in Section 3.

4.1 Reliability

In some contexts, the reliability of defence is more related to the technology than to the relative amount of resources devoted to attack and defence. This section explores a model of defence where the reliability is exogenously specified and we ask how the robust network changes as the technology becomes more reliable.

Suppose there exists a random variable $P : \Omega \to \{0, 1\}$ such that $P(\omega) = 1$ if defence is effective and $P(\omega) = 0$ otherwise. In this Section we shall say that a node is protected if it is defended and defence is effective. Define $p = E(P)$; we refer to p as the index of reliability.

As in the previous section, we shall suppose that defence and attack exhibit indivisibility. If a node i is defended then there is probability p that it is protected. If a node is protected the adversary cannot detect/infect this node, whether directly or indirectly. If, on the other hand, defence is ineffective then the node is vulnerable to direct and indirect detection. Finally, as

before, if a node is detected then infection proceeds along paths leading out of the node and infects all nodes along the path until it encounters a protected node.

Assumption A.2: *Attack and defence decisions are indivisible. A protected node is immune to direct, as well as indirect attack. If a node is unprotected it is vulnerable to both direct and indirect attack. Successful attack on a node leads to elimination of nodes on paths leading away from this node until the process encounters a protected node.*

The general analysis of this problem presents a number of difficulties, and to make progress as well as for easy comparison with earlier results on defence, we restrict attention to a setting in which both the defence and attack budget are set equal to 1. We will focus on the following sequence of moves: $\mathcal{D}$ chooses a network and a node to protect. $\mathcal{A}$ observes this choice and then chooses a node to attack. We shall refer to this game as the DA game with p-defence. Given a strategy for $\mathcal{D}$, let α denote the (only) defended node and C_α its component. Recall, from Proposition 2, that in the basic DA game robust networks consist of 2 equal sized components. With this observation in mind, we restrict attention to networks within $\mathcal{G}_2$.

Proposition 5 *Consider the DA game with p-defence. Suppose* ***A.2*** *holds and let $g \in \mathcal{G}_2$. If $a = 1 = d$ then, in every robust network, C_α is a star. Moreover $\#C_\alpha$ is non-decreasing in p; it equals $n/2$ for p close to 0 and equals n at $p = p_1^* < 1$.*

The proof rests on the following arguments. First we observe that the star with protected center helps the prospects for communication since it minimizes the prospects of indirect node detection. If two components have equal size with one of them protected, $\mathcal{A}$ will find it strictly more attractive to target the non-protected component. By convexity of f, $\mathcal{D}$ strictly improves his payoffs by removing one node from the non-protected component and linking it to the protected center of the other component. This argument can be repeated until $\mathcal{A}$ is indifferent between targeting a node in one or the other component. In other words, the number of nodes in the star component, n_α say, satisfies:

$$pf(n_\alpha - 1) + f(n - n_\alpha) = f(n_\alpha).$$

Since $p \in (0, 1)$ and $f(.)$ is increasing, n_α must be greater than $n/2$. Simple algebra then shows that n_α is increasing in p.

5 Discussion of assumptions

Our analysis yields two main results. One, optimal attack strategies involve targeting a few nodes and ignoring the rest while robust networks consist of equal size groups. Two, if the designer can defend some nodes, robust networks have hub-spoke structure in which the designer protects the hub and the adversary attacks this node. These results were obtained in a model with specific payoff function and with no frictions in communication or detection. We briefly discuss the role of these assumptions now.

5.1 General payoffs

The returns to group size play a key role in our analysis. Recall that, in the basic model, the probability of task completion for a component of size m conditional upon $Q_i = 0$ for all i in that component is given by $f(m) = (\frac{m}{n})^2$. Thus, in particular, $f(.)$ is increasing and convex. We briefly examine more general specifications of the function $f(.)$. First observe that if $f(.)$ is concave or linear, i.e., $f''(.) \leq 0$, then the empty network is optimal, *in the absence of any attack*. The possibility of attack therefore only reinforces this result.

So let us examine robust networks under general convex functions. Suppose the payoffs to the designer from network g under monitoring profile $\mathbf{q}$ are given by

$$V(g, \mathbf{q}) = \sum_{C \in C(g)} f(\#C) \prod_{i \in C} (1 - q_i) \tag{5}$$

where $f(.)$ is increasing and convex.

The following result characterizes equilibrium in the DA game with general convex payoffs.

Proposition 6 *Consider the DA game with payoffs given by (5), and attack budget $a < n$. A robust network consists of k equal size groups, $k \geq a + 1$, and at most one group of smaller size.*

The proof of this result is essentially the same as the proof of Proposition 2 and is omitted. Now consider the effects of changing payoffs on the nature of robust networks. The trade-off with regard to group size is the following: on the one hand, smaller groups lead to more individuals remaining intact after the attack of the adversary. On the other hand, smaller groups exploit fewer of the group size advantage. The group size effect is increasing in convexity of payoffs. To make its effects concrete let us consider the following functional form:

$f(m) = (m/n)^\alpha$, with $\alpha > 1$. Here we interpret α as a measure of the convexity of returns. It is then easy to show that the optimal number of groups is falling, while the size of groups is increasing, in α, the convexity of the payoff function. In general, greater convexity means a greater return from group size, and at the margin this implies fewer and larger groups in the robust network.

We next turn to the analysis of the game with defence. As before, with small defence budgets, the designer has an incentive to design a network with few central nodes and then protect these nodes. The following result is immediate from the proof of Proposition 4.

Proposition 7 *Consider the DA game with defence and payoffs given by (5). Suppose* ***A.1′*** *holds and let $g \in \mathcal{G}_1$. If $d = 1$ then, for arbitrary a, the star is robust.*

We conclude then that our two main findings remain valid if payoffs are increasing and convex in group size.

5.2 Frictions in communication and detection

Frictions in communication and indirect detection are pervasive and present a number of possibilities. A systematic treatment must be left to future work. Here we discuss examples to bring out some of the issues which relate to our main findings.

Suppose that there is a probability $y \in [0, 1]$ such that communication moves across a node to a neighboring node and suppose there is a probability $z \in [0, 1]$ that detection/infection moves across from a node to a neighboring node. Finally suppose, to fix ideas, that the adversary has one unit of attack resources, which she can allocate across the nodes.

First consider the optimality of targeting a single node in a group. This optimality rests on the redundancy involved in attacking two nodes in the same group; this redundancy in turn derives from the frictionless flow of detection/infection across paths of a network. To see this, consider a ring network. Suppose that $z = 0$. Simple computations show that targeting a single node yields payoff $f(n-1)$ to the designer, while equal division of attack across two opposite nodes of a ring network yields $[\frac{1}{2}f(\frac{n}{2}-1)+\frac{1}{4}f(n)+\frac{1}{2}f(n-1)]$. It is easy to check that in the basic model, for which $f(m) = (\frac{m}{n})^2$, equal division of attack dominates targeting a single node for sufficiently large n. In general, the adversary may therefore gain by spreading resources across nodes.

Next consider the robustness of the star network. To maintain consistency with the earlier analysis, suppose that adversary and designer both have one unit of resource ($a = d = 1$) and

that attack and defence decisions are indivisible. Recall, Proposition 3 says that, in the basic model with $z = 1$ and one unit of defence, the robust network is a star with protected center. Now consider the model with $y = 1$ and $z = 0$. It is clear that the equilibrium in the star network again involves the designer protecting the central node and the adversary attacking this node. Payoff to the designer is: $(n-1)f(1)/2 + f(n)/2$. By contrast, the payoff to the designer from a ring network is at least $f(n-1)$. Thus the star with protected center is clearly no longer optimal (irrespective of order of moves).

6 Concluding remarks

This paper studies networks which perform well in the face of attacks.

We first study uniform random attack. Robust networks consist of equal size groups whose number increases in the intensity of the attack. We then turn to intelligent attack: the adversary observes the network and chooses nodes to attack. Optimal attack involves targeting only a few nodes and ignoring the rest. In response, robust networks consist of equal size groups. Their number is higher and performance poorer as compared to the case of uniform random attack.

We extend the framework to allow for defence of nodes. It is attractive to protect central nodes since they minimize the prospects of indirect detection/infection. With limited defence budget it is best to minimize the number of such central nodes: the robust network is a star.

We conclude with a discussion of some open problems. We have throughout assumed that the designer moves first and chooses a network. In some interesting contexts, the adversary may be forced to move first and announce her strategy publicly. The designer then chooses an optimal network given the attack. We have partial results for this adversary move first game: for low attack budgets a robust network is a single group, while for high attack budgets the empty network is robust. Similar results obtain when designer and adversary move simultaneously. The general analysis of these games however remains an open problem.

Our discussion on frictions in communication and detection suggest that they raise novel problems and deserve further research.

Finally, we have focused on the case where the designer chooses the network and the protection strategy; in many applications, such as computer networks, protection is chosen by individual users. The companion paper Goyal and Vigier (2009) studies decentralized linking and protection.

7 Appendix

Proof of proposition 1: (i) and (ii) Notice first that $V((n),0)=1$, while $V(g',0)<1$, $\forall g' \neq (n)$ since in any network with at least two components expert and informer will be disconnected with some probability. It follows by continuity of $V(g,q)$ with respect to q that (n) is uniquely optimal on some neighborhood of zero. We now proceed to characterize the maximal size of this neighborhood.

Recall that $g^{b2}=(\frac{n}{2},\frac{n}{2})$, and let $\widetilde{q}_n$ denote the monitoring intensity such that $V((n),\widetilde{q}_n)=V(g^{b2},\widetilde{q}_n)$. Expanding V gives

$$(1-\widetilde{q}_n)^n=\frac{1}{2}(1-\widetilde{q}_n)^{n/2}$$

and, solving for $\widetilde{q}_n$,

$$\widetilde{q}_n=1-(\frac{1}{2})^{2/n} \tag{6}$$

Notice that $V((n),q)>V(g^{b2},q)$ to the left of $\widetilde{q}_n$, while the reverse inequality holds to the right of $\widetilde{q}_n$.

Step 1: We show first that, for $q=\widetilde{q}_n$, $(\frac{n}{2},\frac{n}{2})$ is uniquely optimal within $\mathcal{G}_2\backslash\mathcal{G}_1$.

Fix $q=\widetilde{q}_n$ and $g\in\mathcal{G}_2$ with components size $\frac{n}{2}(1-x)$ and $\frac{n}{2}(1+x)$, $x\in[0,1)$. We have,

$$V(g,\widetilde{q}_n)=\left(\frac{\frac{n}{2}(1-x)}{n}\right)^2(1-\widetilde{q}_n)^{\frac{n}{2}(1-x)}+\left(\frac{\frac{n}{2}(1+x)}{n}\right)^2(1-\widetilde{q}_n)^{\frac{n}{2}(1+x)}$$

which using (6) gives

$$V(g,\widetilde{q}_n)=\frac{1}{4}(1-x)^2(\frac{1}{2})^{(1-x)}+\frac{1}{4}(1+x)^2(\frac{1}{2})^{(1+x)} \tag{7}$$

Simple calculus then shows that (7) has a unique maximum for $x\in[0,1)$, attained at $x=0$, i.e. when both components have equal size.

Step 2: We show that $V((n),q)>V(g,q)$, $\forall\ q<\widetilde{q}_n$, $\forall g\in\mathcal{G}_2\backslash\mathcal{G}_1$.

Let $g\in G_2\backslash(n)$ with components size a and a, and $q<\widetilde{q}_n$. We have

$$V((n),q)-V(g,q)=(1-q)^n-\left[\left(\frac{a}{n}\right)^2(1-q)^a+\left(\frac{b}{n}\right)^2(1-q)^b\right] \tag{8}$$

Recall that $V((n),0)-V(g,0)>0$. Next, note that $V((n),\widetilde{q}_n)-V(g,\widetilde{q}_n)=V(g^{b2},\widetilde{q}_n)-$

$V(g, \widetilde{q}_n) > 0$ following the definition of $\widetilde{q}_n$ and using step 1. Since by Descartes' sign rule[4] for the polynomial $X^n - \left(\frac{a}{n}\right)^2 X^a - \left(\frac{b}{n}\right)^2 X^b$ we know that (8) crosses the intercept at most once on $q \in [0, 1)$ we conclude that $V((n), q) - V(g, q) > 0, \forall q < \widetilde{q}_n$ as indicated.

Step 3: We show that for $q = \widetilde{q}_n$ no optimal network can have two components of unequal size.

In what follows, we use the following notation: given $g \in \mathcal{G}$, and $C \in C(g)$, we let V_C denote the probability of task completion in component C.

Fix $q = \widetilde{q}_n$ and consider $g \in \mathcal{G}$ with two components of unequal size. If these are g's only components then by step 1 we know that g can be improved upon by taking balanced components. Assume therefore that g has three components at least, and let C_1, $C_2 \in C(g)$, $\#C_1 \neq \#C_2$. Let $m = \#C_1 + \#C_2$. By component additivity it is enough to show that $V_{C_1 \cup C_2}$ can be improved upon. Using (6) note that $\frac{d\widetilde{q}_n}{dn} < 0$, and so $\widetilde{q}_m > \widetilde{q}_n$ since by assumption $m < n$. But then step 2 shows that (m) dominates $(\#C_1, \#C_2)$ at $q = \widetilde{q}_n$. This completes step 3.

Step 4: We show that, for $q = \widetilde{q}_n$, $(\frac{n}{2}, \frac{n}{2})$ is uniquely optimal within $\mathcal{G} \backslash \mathcal{G}_1$.

By step 3, for $q = \widetilde{q}_n$ all unbalanced networks can be ignored. It therefore only remains to check that $(\frac{n}{2}, \frac{n}{2})$ is optimal within the family of balanced networks, $\{(\frac{n}{k}, .., \frac{n}{k}), k \geq 2\}$. We have

$$V(g^{bk}, \widetilde{q}_n) = k \left(\frac{1}{k}\right)^2 (1 - \widetilde{q}_n)^{\frac{n}{k}}$$

and, substituting for $\widetilde{q}_n$ using (6),

$$V(g^{bk}, \widetilde{q}_n) = \frac{1}{k} (\frac{1}{2})^{\frac{2}{k}}$$

To complete, notice that $\frac{\partial V(g^{bk}, \widetilde{q}_n)}{\partial k} = k^{-2} (\frac{1}{2})^{\frac{2}{k}} \left[\frac{2 \ln 2}{k} - 1\right] < 0$, $\forall k \geq 2$.

We may now complete the proof of (i) and (ii). Recall that (n) is uniquely optimal for $q = 0$. Since by step 4 $(\frac{n}{2}, \frac{n}{2})$ is optimal for $q = \widetilde{q}_n$ and $V((n), \widetilde{q}_n) = V(g^{b2}, \widetilde{q}_n)$, (n) is also optimal for $q = \widetilde{q}_n$. It follows from Descartes' sign rule that (n) is optimal over the interval $[0, \widetilde{q}_n]$, uniquely except at $\widetilde{q}_n$. Point (ii) follows from step 4 in the proof and continuity of V with respect to q.

(iii) Let $g \in \mathcal{G}$ and let $C \in C(g)$ denote a component of size $m \geq 2$. Consider the value of connecting the m^{th} node to C as opposed to leaving that node isolated. This is given

[4]Descartes' sign rule may be stated as follows: When a polynomial function is written in standard form, the number of changes in sign of the coefficients is the maximum number of positive zeros of the function.

by

$$2(\frac{1}{m})\left(\frac{m-1}{m}\right)(1-q)^m - \left[\left(\frac{m-1}{m}\right)^2 q(1-q)^{m-1} + \left(\frac{1}{m}\right)^2 (1-q)(1-\left(1-q\right)^{m-1})\right]$$

Substituting $q = \frac{1}{2}$ and rearranging gives,

$$\frac{1}{m^2}(\frac{1}{2})^{m-1}\left[(m-1)\left(\frac{3-m}{2}\right) - \frac{1}{2}\frac{1-(\frac{1}{2})^{m-1}}{(\frac{1}{2})^{m-1}}\right]$$

Simple calculus shows that this expression is zero when $m = 2$ and strictly negative for all $m > 2$. Therefore, the balanced network made up of paired nodes is optimal for $q = \frac{1}{2}$, while any network containing a component of size ≥ 3 can be strictly improved upon by disconnecting one node in that component.

By continuity of V with respect to q, there exists $\zeta_n > 0$ such that $(2, .., 2)$ strictly dominates any network containing a component of size ≥ 3 for $q \in (\frac{1}{2} - \zeta_n, \frac{1}{2}]$.

Since the value of connecting two isolated nodes is given by $(1-q)^2 - (\frac{1}{2})(1-q) = (1-q)(\frac{1}{2} - q) > 0$, $\forall q < \frac{1}{2}$, we also have $V(g^{b\frac{n}{2}}, q) > V(g^e, q)$, $\forall q < \frac{1}{2}$.

It therefore follows that $(2, ..2)$ is uniquely optimal for $q \in (\frac{1}{2} - \zeta_n, \frac{1}{2})$, as indicated in (iii).

(iv) We proceed in two steps, which allow us to make slightly stronger statements.

Step 1: No agent with $q_i > \frac{1}{2}$ may belong to a component of size $m > 2$ in equilibrium.

Let $g \in \mathcal{G}$, $C \in C(g)$, $\#C = m > 2$. Consider the value of i belonging to C as opposed to leaving i isolated. This is given by

$$\begin{aligned}
& 2\left(\frac{1}{m}\right)\left(\frac{m-1}{m}\right)(1-q_i)\prod_{j\in C\setminus\{i\}}(1-q_j) - \left(\frac{m-1}{m}\right)^2 q_i \prod_{j\in C\setminus\{i\}}(1-q_j) \\
& -(\frac{1}{m})^2(1-q_i)(1-\prod_{j\in C\setminus\{i\}}(1-q_j)) \\
\leq\ & 2\left(\frac{1}{m}\right)\left(\frac{m-1}{m}\right)(1-q_i)\prod_{j\in C\setminus\{i\}}(1-q_j) - \left(\frac{m-1}{m}\right)^2 q_i \prod_{j\in C\setminus\{i\}}(1-q_j) \\
=\ & [2\left(\frac{1}{m}\right)\left(\frac{m-1}{m}\right)(1-q_i)\prod_{j\in C\setminus\{i\}}(1-q_j)](1-\frac{(m-1)}{2}\frac{q_i}{(1-q_i)}) \\
<\ & 0 \qquad \forall m \geq 3
\end{aligned}$$

Step 2: No two agents with monitoring intensity strictly more than $\frac{1}{2}$ may be connected

in equilibrium.

Let $i,\ j \in N,\ i \neq j,\ q_i > \frac{1}{2}$, and $q_j > \frac{1}{2}$. By step 1, both i and j must belong to a component of size ≤ 2 in equilibrium. It is enough therefore to check that $\{i,j\}$ cannot form a component of size 2 in equilibrium. Assume without loss of generality that $q_i \leq q_j$. The value of link ij is given by

$$\begin{aligned}
& \left(\frac{1}{2}\right)(1-q_i)(1-q_j) - \left(\frac{1}{2}\right)^2 q_i(1-q_j) - \left(\frac{1}{2}\right)^2 q_j(1-q_i) \\
\leq\ & \left(\frac{1}{2}\right)(1-q_i)(1-q_j) - \left(\frac{1}{2}\right) q_i(1-q_j) \\
<\ & \left(\frac{1}{2}\right)(1-q_i)(1-q_j) - \left(\frac{1}{2}\right)(1-q_i)(1-q_j) \\
=\ & 0
\end{aligned}$$

Point (iv) then follows immediately from step 2.

(v) Consider $g \in \mathcal{G}$, with $C_1,\ C_2 \in C(g)$, $\#C_1 > \#C_2 = a$. Let $k > 1$ s.t. $\#C_1 = ka$. We claim that g is never optimal, $\forall q \in [0,1]$. By component additivity it is enough to show that $V_{C_1 \cup C_2}$ can be improved upon, $\forall q \in [0,1]$. Let $m = (k+1)a$. We want to show that (ka, a) is never optimal within the set of networks on m nodes. Using Descartes' sign rule we know that $\exists q_\alpha,\ q_\beta \in (0,1)$ s.t. (ka,a) dominates $(\frac{m}{2}, \frac{m}{2})$ on $[0, q_\alpha) \cup (q_\beta, 1]$ while $(\frac{m}{2}, \frac{m}{2})$ dominates on (q_α, q_β). Moreover, we know by (i) and (ii) that (ka,a) is dominated by (m) on $[0, q_\alpha)$. We can conclude that (ka,a) is never optimal on $[0, q_\beta)$. Our next step consists in showing that $(a, a, .., a)$ dominates (ka,a) on $(q_\beta, 1]$. Notice that $V((a,a,..,a),q) = (k+1)V_{C_2}$. Since $V((ka,a),q) = V_{C_1} + V_{C_2}$, this is equivalent to showing that $kV_{C_2} > V_{C_1}$ on $(q_\beta, 1]$. It is easy to show that $\exists! q^\dagger$ such that $V_{C_1} > kV_{C_2}$ on $[0, q^\dagger)$ while $kV_{C_2} > V_{C_1}$ on $(q^\dagger, 1]$. Therefore, it is enough to show that $q^\dagger < q_\beta$. Next, by definition of q_β, it is enough to show that $(\frac{m}{2}, \frac{m}{2})$ dominates (ka,a) at $q = q^\dagger$. Simple algebra shows that $q^\dagger = 1 - (\frac{1}{k})^{\frac{1}{(k-1)a}}$, giving

$$V((ka,a), q^\dagger) = (\frac{k}{k+1})^2 (\frac{1}{k})^{\frac{k}{k-1}} + (\frac{1}{k+1})^2 (\frac{1}{k})^{\frac{1}{k-1}}$$

and

$$V((\frac{m}{2}, \frac{m}{2}), q^\dagger) = \frac{1}{2}(\frac{1}{k})^{\frac{k+1}{2(k-1)}}$$

Simple calculus then shows that $V((\frac{m}{2}, \frac{m}{2}), q^\dagger) > V((ka,a), q^\dagger)$, $\forall k > 1$.

■

Proof of Proposition 2: For $a = n$, $\mathcal{A}$ monitors all nodes with intensity one and the task is never carried out. Suppose $a < n$. The proof proceeds in four steps.

Step 1: First, notice that by redundancy it is always strictly suboptimal for $\mathcal{A}$ to monitor two nodes with strictly positive intensity within a single component. For any fixed network, $\mathcal{A}$ therefore monitors at most one node in any component in equilibrium.

Note in particular that by Step 1, in equilibrium and given any network, $\mathcal{A}$ allocates one unit at most of her budget to any component. Henceforth, we say that $\mathcal{A}$ removes a component if it monitors a single node with intensity 1 in it.

Step 2: Let $\overline{s}(g)$ denote the maximal component size in g. We show that in equilibrium $a+1$ components at least have size $\overline{s}(g)$. Suppose this is not the case. Let C_1 have maximal size in g. Then, in equilibrium, C_1 is removed. If this were not the case, using Step 1 and the fact that at most a components have maximal size in g, we could find a strictly smaller component in which some node is monitored with positive intensity. But then $\mathcal{A}$ strictly improves her payoffs by shifting ε from that node to C_1. Next, consider forming g' from g in which C_1' is obtained from C_1 by isolating a single node, leaving the rest of the network unchanged. In g', either C_1' has maximal size, or at most $a-1$ components have size strictly greater than it. Hence, without loss of generality in payoffs, we may assume that C_1' is removed in equilibrium. But then $\mathcal{D}$ does strictly better with g' than g since by doing so she saves the node she isolated. This contradicts the equilibrium assumption on g, and concludes Step 2.

Step 3: We show that in equilibrium at most one component has size less than $\overline{s}$. Let g denote an equilibrium network. By Step 2, $a+1$ components at least have maximal size in g. Thus if two components have less than maximal size, neither can be monitored in equilibrium. But then notice that by repeated use of Step 2, and convexity of returns to size, $\mathcal{D}$ strictly improves her payoffs by shifting one node from the smaller to the larger of the two components.

Step 4: We show that if $n \text{mod} 2a = 0$, then $\mathcal{D}$ chooses $2a$ components of equal size. First we show that a network with $2a$ components of equal size is optimal within the class of balanced networks (i.e. such that all components have equal size). Let s denote the equal size of all components. Equilibrium payoffs to $\mathcal{D}$ are given by $\left(\frac{s}{n}\right)^2 \left(\frac{n}{s} - a\right)$. This is quadratic in s with a unique maximum attained for $s = \frac{n}{2a}$. So $2a$ components are uniquely optimal within the class of balanced networks. By Step 3, we only have to check that no network does better with a single component having less than maximal size. Let g denote such a network, with β components of maximal size $\overline{s}$ (by Step 2, assume without loss of generality $\beta \geq a+1$) and

one component of size $s < \overline{s}$. Equilibrium payoffs to $\mathcal{D}$ are then

$$\left(\frac{\overline{s}}{n}\right)^2 (\beta - a) + \left(\frac{s}{n}\right)^2 \tag{9}$$

Now, using $n = \beta\overline{s} + s$, notice that

$$\begin{aligned} \left(\frac{\overline{s}}{n}\right)^2 (\frac{n}{\overline{s}} - a) &= \left(\frac{\overline{s}}{n}\right)^2 (\beta + \frac{s}{\overline{s}} - a) \\ &= \left(\frac{\overline{s}}{n}\right)^2 (\beta - a) + \frac{\overline{s}s}{n^2} \end{aligned} \tag{10}$$

and (9)<(10) since $s < \overline{s}$. But we have also shown that (10)< $(\frac{1}{2a})^2(2a - a) = \frac{1}{4a}$, hence in equilibrium $\mathcal{D}$ does strictly better with $2a$ components of equal size and the proof is complete.

For $a > \frac{n}{2}$, just note that $\mathcal{A}$ has enough budget that in equilibrium any component of size 2 or more is removed. The empty network is therefore the unique equilibrium outcome. ■

The next Lemma is useful in the proofs of the subsequent results.

Lemma 1 *For any connected network g there exists a node i with the property that, for any $j \in g$ fixed, $j \neq i$, there exist j-independent paths connecting i and $\frac{n}{2}$ nodes in g at least.*

Proof: We work the proof for minimally connected networks. If g is not minimally connected then there exists a minimally connected network g' obtained from g by deletion of links. Then a node i satisfying the property in g' also satisfies it in g.

The proof is by induction on n, the total number of nodes. For $n = 2$ the property is obviously satisfied. Let $n > 2$ and assume the property holds for any network with $n - 1$ or less nodes. Let g be minimally connected with n nodes. Consider g' obtained from g by removing a leaf l in g (i.e. a node with degree 1). Using the induction hypothesis on g' we can find i' satisfying the property for g'. Next, let i denote the neighbour of i' on the unique path between i' and l in g. We show that one of i or i' must satisfy the property for g. If i' satisfies the property for g then we are done. Suppose i' fails to satisfy the property for g. Let $Y_{s\backslash t}(g)$ ($y_{s\backslash t}(g)$) denote the set (cardinality) of nodes which can be connected to s in g through some t-independent path. Note that $Y_{i\backslash i'}(g) = N\backslash Y_{i'\backslash i}(g)$, and so $y_{i\backslash i'}(g) = n - y_{i'\backslash i}(g)$. Since $y_{i'\backslash i}(g) < \frac{n}{2}$ by hypothesis, it follows that $y_{i\backslash i'}(g) > \frac{n}{2}$. Next, let j denote a neighbour of i in g other than i'. Note that $y_{i\backslash j}(g) \geq y_{i'\backslash i}(g')+1$. By definition of i' we have $y_{i'\backslash i}(g') \geq \frac{n-1}{2}$. Hence

$y_{i\setminus j}(g) \geq \frac{n-1}{2} + 1 > \frac{n}{2}$. Thus we have shown that for any neighbour t of i in g, $y_{i\setminus t}(g) \geq \frac{n}{2}$. Since for any node non-neighbour t' of i there exists a neighbour t with $y_{i\setminus t'}(g) > y_{i\setminus t}(g)$, the proof is complete. ■

Proof of Proposition 3: The proof is in two steps. We first show that the star is robust within $\mathcal{G}_1$. We then turn to the optimality of a single component.

Fix a network g with one component. We show that, for any order of moves, $\mathcal{A}$ can guarantee herself payoff $-\frac{1}{2}f(n)$. By Lemma 1, we can find a node i which can access half the nodes at least when any one node other than itself is protected. Now consider the designer move first game. For any network with one component, fix some defended node j. The adversary can ensure herself a payoff $f(n)/2$ by targeting node j. Next consider the adversary move first game: suppose the adversary targets node i. If $\mathcal{D}$ defends node i she receives payoff $\frac{1}{2}f(n)$. If she defends some other node, then by definition of node i she receives a maximum of $f(n/2)$. By convexity, $f(\frac{n}{2}) < \frac{1}{2}f(n)$, and so $\mathcal{A}$ can guarantee herself payoff $-\frac{1}{2}f(n)$. Finally, note that in the simultaneous move game, the adversary can ensure herself a minimum of $f(n)/2$ by picking a node i, identified in Lemma 1.

Next, it is easy to see that given a single star, and for any order of moves, the unique equilibrium involves $\mathcal{A}$ and $\mathcal{D}$ both targeting the center. Since the resulting payoff to $\mathcal{A}$ is $-\frac{1}{2}f(n)$, $\mathcal{D}$ can do no better and so the star is robust within $\mathcal{G}_1$.

In order to show that the star is robust more generally when more components are allowed, we obtain maximum payoffs for networks with 2 or more components and show that they are lower than the payoffs that designer earns in the single star with protected center strategy.

We first provide the proof for $g \in \mathcal{G}_2$. From earlier arguments, we may restrict attention to networks in $\mathcal{G}_2$ in which all components are stars. Details of the analysis vary with the order of moves considered, and so we analyze each case separately.

$\mathcal{A}$ moves first: Consider two stars with size $\frac{n}{2} + x$ and $\frac{n}{2} - x$ respectively. The unique equilibrium in this subgame involves $\mathcal{A}$ and $\mathcal{D}$ both targeting the central node in the largest component. Payoff to $\mathcal{D}$ is thus $\frac{1}{2}f(\frac{n}{2} + x) + f(\frac{n}{2} - x)$. By convexity of f the maximum is thus attained at $x = 0$ or $x = \frac{n}{2}$. It is easily checked that payoffs for $x = 0$ are given by $\frac{3}{8}$. The single star is thus robust in this case.

$\mathcal{D}$ moves first: As before, consider two stars with size $\frac{n}{2}+x$, and $\frac{n}{2}-x$ respectively. Any equilibrium in this subgame involves $\mathcal{D}$ defending the central node in the largest component. $\mathcal{A}$ then either chooses to attack the same node or some node in the smallest component, whichever gives her highest payoff. Equilibrium payoffs to $\mathcal{D}$ in this subgame are thus given

by

$$\min\{\frac{1}{2}f(\frac{n}{2}+x)+f(\frac{n}{2}-x), f(\frac{n}{2}+x)\} \tag{11}$$

Since we have shown earlier that $\max_{0\leq x\leq \frac{n}{2}}\{\frac{1}{2}f(\frac{n}{2}+x)+f(\frac{n}{2}-x)\}=\frac{1}{2}$, the single star is also robust in this case.

Simultaneous moves: Again, consider two stars with size $\frac{n}{2}+x$, and $\frac{n}{2}-x$ respectively. For x sufficiently large that $\frac{1}{2}f(\frac{n}{2}+x)\geq f(\frac{n}{2}-x)$, the equilibrium in the subgame is in pure strategies, with both players targeting the central node in the largest component. The analysis in this case can thus be borrowed from previous cases. For x small however, so that $\frac{1}{2}f(\frac{n}{2}+x)<f(\frac{n}{2}-x)$, the equilibrium involves both players mixing between targeting the central node of the largest component and targeting the central node of the smaller component. Let $q_{\mathcal{A}}$ denote the equilibrium probability with which $\mathcal{A}$ targets the central node of the largest component. The indifference condition for $\mathcal{D}$ is

$$q_{\mathcal{A}}[\frac{1}{2}f(\frac{n}{2}+x)+f(\frac{n}{2}-x)]+(1-q_{\mathcal{A}})f(\frac{n}{2}+x)=(1-q_{\mathcal{A}})[f(\frac{n}{2}+x)+\frac{1}{2}f(\frac{n}{2}-x)]+q_{\mathcal{A}}f(\frac{n}{2}-x)$$

Let $f^+=f(\frac{n}{2}+x)$, and $f^-=f(\frac{n}{2}-x)$. We obtain

$$q_{\mathcal{A}}=\frac{f^-}{f^++f^-}$$

And equilibrium payoffs to $\mathcal{D}$ are given by

$$\frac{f^-}{f^++f^-}[\frac{1}{2}f^++f^-]+\frac{f^+}{f^++f^-}f^+ \tag{12}$$

It is easily checked that payoffs given in (12) are less than $\frac{1}{2}$ as x ranges from 0 to x^*, where x^* is obtained by solving the condition $\frac{1}{2}f(\frac{n}{2}+x)=f(\frac{n}{2}-x)$, i.e. $x^*=n\frac{\sqrt{2}-1}{2(\sqrt{2}+1)}$. Thus, a single star is robust in the simultaneous moves case too.

■

Proof of Proposition 4: Take the star network with protected center and let α denote its central node. Let t denote the units of attack allocated to α in the adversary's optimal response. Next, consider an arbitrary connected network g' in which a single (arbitrary) node, α' say, is defended. Consider the following attack strategy: allocate t units of attack to node α', and 1 unit to $a-t$ other nodes. It is easy to see that by following this strategy for all networks the adversary guarantees herself payoffs at least as high as the maximum payoffs

she receives when the network is a star with protected center. Thus, the star with protected center is robust within the class of connected networks.

■

Proof of Proposition 5: In this proof, with a slight abuse of notation, we let g denote a strategy for the designer. Thus, g effectively consists of a network and a defended node in it.

First, note that choosing C_α to be a star centered at α is un-dominated for $\mathcal{D}$. To see this, start with a design g such that this is not the case and consider reordering nodes in C_α to form a star centered at α leaving the rest of the network unchanged. Let g' denote the resulting design. Clearly

$$V(g', \mathbf{q}) \geq V(g, \mathbf{q}) \qquad \forall \mathbf{q} \in [0,1]^n$$

Hence we can find a network that is robust and in which C_α is a star centered at α.

Next, we characterize the equilibria in which C_α is a star centered at α. We say that $\mathcal{D}$ plays n_α if he chooses a network in which C_α forms a star with n_α nodes centered at α. In turn, a best-response for $\mathcal{A}$ solves:

$$\begin{aligned}
\min_{q_\alpha+q=1} \pi^{\mathcal{D}} &= \min_{q_\alpha+q=1} p[(1-q_\alpha)f(n_\alpha) + q_\alpha f(n_\alpha - 1)] + (1-p)(1-q_\alpha)f(n_\alpha) + (1-q)f(n-n_\alpha) \\
&= \min_{q_\alpha+q=1} (1-q_\alpha)f(n_\alpha) + q_\alpha p f(n_\alpha - 1) + (1-q)f(n-n_\alpha) \\
&= \min_{q_\alpha+q=1} f(n_\alpha) + f(n-n_\alpha) - q_\alpha[f(n_\alpha) - pf(n_\alpha - 1)] - qf(n-n_\alpha) \qquad (13)
\end{aligned}$$

where $q_\alpha \in \{0,1\}$ indicates the monitoring intensity of a peripheral node chosen in C_α, and $q \in \{0,1\}$ the monitoring intensity of a node chosen in $N \backslash C_\alpha$. Inspection of (13) gives $q_\alpha = 1$ if $f(n_\alpha) - pf(n_\alpha - 1) > f(n - n_\alpha)$, and $q = 1$ otherwise. Next, notice that $f(n_\alpha) - pf(n_\alpha - 1)$ increases in n_α by convexity of f, while $f(n - n_\alpha)$ trivially decreases in n_α. The two curves therefore cross at most once, and $\exists!\ n^*_\alpha(p) \in [0,1]$ such that in equilibrium $\mathcal{A}$ targets a single node with intensity one in C_α if $n_\alpha > n^*_\alpha(p)$, and a single node with intensity one in C if $n_\alpha < n^*_\alpha(p)$. In turn we deduce that equilibrium payoffs to $\mathcal{D}$ from playing n_α are given by

$$\begin{aligned}
h_1(n_\alpha; p) &= pf(n_\alpha - 1) + f(n - n_\alpha)\ , \quad n_\alpha > n^*_\alpha(p) \qquad (14) \\
h_2(n_\alpha) &= f(n_\alpha)\ , \quad n_\alpha < n^*_\alpha(p).
\end{aligned}$$

and, by definition, $h_1 = h_2$ at $n^*_\alpha(p)$.

Next, we establish the following points: (i) $n_\alpha \geq n^*_\alpha(p)$ (in particular, it is always a best-response for $\mathcal{A}$ to target a single node with intensity one in C_α); (ii) if $n_\alpha > n^*_\alpha(p)$ then $n_\alpha = n$; (iii) in equilibrium, n_α is non-decreasing in p.

Point (i) follows by noting that f is increasing. Point (ii) follows from convexity of h_1 in its first argument (which is a consequence of the convexity of f).

We establish point (iii) in two steps. First, we show that $\exists\ p^*_1 \in [0,1]$ such that $n_\alpha = n^*_\alpha(p)$ is robust for $p < p^*_1$, while $n_\alpha = n$ is robust for $p > p^*_1$. This is true since $\frac{\partial}{\partial n_\alpha}(\frac{\partial h_1}{\partial p}) = f'(n_\alpha - 1)$, and $f'(n_\alpha - 1) > 0$ in equilibrium (note that no network is robust in which $n_\alpha < \frac{n}{2}$, else just switch α with a node in $N \backslash C_\alpha$ to find a strict improvement on $\mathcal{D}$'s strategy in equilibrium). Now point (iii) is immediate if $p \geq p^*_1$, and if $p < p^*_1$ the result follows from observing that $\frac{\partial h_1}{\partial p} > 0$ and $\frac{dh_2}{dn_\alpha} > 0$. Also, we can see from (14) that $n_\alpha = \frac{n}{2}$ is robust for p close to zero, and that $n_\alpha = n$ is robust for p close to one. In particular, $0 < p^*_1 < 1$.

Lastly, we show that a network with C_α a star centered at α is uniquely robust. Let g^+ denote a design in which C_α isn't a star centered at α. If $\#C_\alpha < \frac{n}{2}$, g is strictly dominated in equilibrium as noted previously. Suppose then $\#C_\alpha \geq \frac{n}{2}$. Two unprotected nodes at least are adjacent in C_α. We next distinguish two cases: (i) If a best-response for $\mathcal{A}$ involves monitoring a node in C_α then clearly $\mathcal{D}$ does strictly better by reorganizing C_α to form a star centered at α; (ii) If monitoring a node in C_α is not a best-response for $\mathcal{A}$, then payoffs are unchanged if $\mathcal{D}$ plays $n_\alpha = \#C_\alpha$ (i.e., plays the strategy in which links in C_α are reorganized to form a star centered at α) and moreover, in that sub-game, it is still not a best-response for $\mathcal{A}$ to monitor a node in C_α.

■

8 References

1. Albert R, Jeong H, Barabási, A-L (2000), Error and attack tolerance of complex networks, *Nature*, 406: 378-82.

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

3. Barabasi, A-L (1999), *Linked*. Perseus Books.

4. Bala, V. and Goyal, S. (2000), An analysis of strategic reliability, *Review of Economic Design*, 5, 205-28.

5. Bolton, P. and M. Dewatripont (1994), The firm as a communication network, *Quarterly Journal of Economics*, 109, 809-839.

6. Farley, J. D. (2003), Breaking Al Queda Cells: A mathematical analysis of counter-terrorism operations, *Studies in Conflict and Terrorism*, 26, 399-411.

7. Farley, J. D. (2006), Building the perfect terrorist cell, Conference Talk.

8. Garicano, L. (2000), Hierarchies and the Organization of Knowledge in Production, *Journal of Political Economy*, volume 108, pages 874-904.

9. Garoupa, N. (2007), Optimal Law enforcement and criminal organization, *Journal of Economic Behavior and Organization*, 63, 461-474.

10. Goyal, S. and A. Vigier (2009), Interaction, infection and control. *Mimeo*, Cambridge University.

11. Goyal, S. (2007), *Connections: an introduction to the economics of networks.* Princeton University Press.

12. Grotschel, M., C.L. Monma and M. Stoer (1995), Design of survivable communication networks, in M.O. Ball, TL. Magnanti, C.L. Monma and G.L. Nemhauser (eds) *Handbooks of Operations Research and management science: Network Models.* North Holland. Amsterdam, 617-672.

13. Hong, S. (2008), Hacking-proofness and Stability in a Model of Information Security Networks, working paper.

14. Krueger, A. (1974), The Political Economy of the Rent-Seeking Society, *American Economic Review* 64, 3, 291303.

15. Levine, S. (1999), *Fragile Dominion: Complexity and the Commons* Perseus Books, Reading, MA.

16. Nagaraja, S., Anderson, R. (2007) The topology of covert conflict, *Cambridge Computer Laboratory Technical Report 637.*

17. Radner, R (1992), Hierarchy: The Economics of Managing, *Journal of Economic Perspectives*, 30, 3, 1382-1415.

18. Radner, R. (1993), The organization of decentralized information processing, *Econometrica*, 61, 5, 1109-1146.

19. Smith, C. J (2008), Preface to special issue on *Networks: Games, Interdiction, and human interaction problems on networks*, Volume 52, 3, 109-110.

20. Tullock, G. (1967), The Welfare Costs of Tariffs, Monopolies, and Theft, *Western Economic Journal* 5, 3, 224232.

21. Van Zandt, T. (1999), Decentralized information processing in the theory of organizations, *Contemporary Economic Issues Volume 4: economic design and behavior*, edited by Murat Sertel. MacMillan Press. London.