OFR BRIEF SERIES
OFFICE OF FINANCIAL RESEARCH

18-02 | August 1, 2018

Network Analysis: Defending Financial Stability by Design

By Stacey Schreft and Simpson Zhang[1]

The financial system operates through a complex set of networks, and a variety of operational failures can disrupt those networks. If defenses and recovery efforts fail, such disruptions can threaten financial firms or even the stability of the financial system. This brief explains network modeling, one tool the OFR uses to understand threats to financial stability. Models reveal the most important players in networks and the most effective defense strategies. Some network structures are more resilient to random failures, such as those caused by natural disasters. Others are more resilient to targeted incidents, such as hacks. Coordination among regulators and industry is vital for a strong defense.

The modern financial system relies on a web of networks: electrical networks such as power grids, technical networks such as computer systems, and financial networks such as intermediation chains. Each network is exposed to its own strain of threats and needs its own self-protection strategy. However, networks are also dependent on other networks. Operational failures or incidents in one network can spread through contagion to the rest of the system.

Consider the power outage that began on Aug. 14, 2003. It started with a high-voltage line failing in Ohio. The electric power grid through which utilities share power when needed allowed the outage to cascade through southern Canada and the northeastern United States, leaving as many as 50 million people in the dark, some for days. All told, the event caused about $6 billion in losses.[2] What started as an electrical outage ended up having far-ranging impacts.

Another example is the September 2017 hack of Equifax, the credit-reporting agency. Hackers reportedly stole financial information about 145 million Americans.[3] The hackers' goal might have been only the theft of information that could be resold or used for other purposes, not to disable Equifax or threaten the financial system more broadly. Nevertheless, the damage to the financial system could have been greater. Operational incidents can threaten financial stability through three main channels. They can (1) disrupt the operations of a financial firm that provides critical services, (2) reduce confidence in firms and markets, or (3) damage the integrity of key data. The OFR has highlighted these contagion channels in previous

publications.[4] Any of those three channels could provide a path for contagion from one firm to the wider financial system.

At the OFR, we study networks to better understand the plumbing of the financial system and identify potential weaknesses in network structures and defenses. Economists use a number of tools to study different types of networks, including electrical networks, technical networks, or financial intermediation networks. These "network models" can help regulators and the financial industry better understand the security and resilience of financial networks and thus assist them in designing better defenses to reduce the three channels of risk discussed above.

The OFR has been analyzing networks since its inception. Much of that analysis has focused on the exposures of financial firms to counterparties through the firms' asset holdings. One firm's default can set off a cascade of defaults to other firms through the network of exposures. For instance, an OFR working paper on the market for credit default swaps traded through central counterparties (CCPs) showed that peripheral firms may be a greater source of contagion than the CCP itself.[5]

The network models discussed in this brief extend our previous work on contagion to a range of different markets and network types.

Network models also bring to light real-world design trade-offs. For example, models reveal that some network defenses are more resilient to random failures, such as from natural disasters or random power outages. Other defense configurations are more resilient to targeted incidents, such as computer-system hacks.

This brief shows how network models help in understanding the impacts of operational incidents and in building better defenses. It then looks at the dangers of decentralized approaches to network security and recommends other approaches.

Network Model Basics

Networks are fundamental building blocks of the financial system. Most major financial markets operate through networks — not only information technology networks, but also other kinds of relationships. Credit default swaps trading, interbank lending, and equity markets operate through networks, some more formal than others. The operational, messaging, and computer systems underlying these markets can be represented as networks. Formal and informal networks are potential sources of financial instability because financial or operational failures could spread through them. Understanding how networks operate can help promote financial stability.

Although financial and computer networks are complex, they can be represented in a simple fashion through the framework of *network analysis*.

Network analysis starts by modeling a network as a group of entities represented by *nodes* and *links.* Each node represents a participant. Depending on the type of network, a participant could be a financial institution, a component of a firm's computer system, or a payment hub.

A link between two nodes represents a connection between two participants. This connection could represent a variety of interactions depending on the market or system being modeled. For example, it could represent a financial exchange, such as an extension of credit, or a technical interaction, such as messaging between two parties. A link could even represent shared exposures, as might arise for firms using the same software platform.

Two linked nodes are called *neighbors.* A *path* between two nodes exists if a series of links goes from one node to the other. The *shortest path* is the one with the fewest links between nodes.

An *adjacency matrix* is a table of ones and zeroes that helps in analyzing a network. A one in the matrix represents a link between two nodes. A zero represents no link. **Figure 1** shows a simple network with its adjacency matrix. Node B is linked to Nodes A and C, while Nodes A and C are not linked directly. However, a path connects Nodes A and C through Node B. The length of this path can be described as two because two links must be traversed.

How Important Is Each Part of a Network?

An adjacency matrix enables calculation of several network metrics. Centrality metrics identify the significance of a node (for example, a market participant or firm) to a network. This information helps pinpoint the nodes most vulnerable to incidents or most critical to protect.

Different centrality metrics are useful for different purposes. For instance:

- **Degree centrality** measures the number of neighbors of a node. The more neighbors a node has, the more entities that node interacts with and so the more critical the node is to the network.
- **Betweenness centrality** measures how often a node is in the shortest path connecting two other nodes. A node in more of the shortest paths means the node could be involved in intermediating more transactions with other nodes. Its disruption would impair these transactions.
- **Closeness centrality** measures how close on average a node is to other nodes through the shortest path. Being closer to other nodes means the node can interact with those nodes more easily, or spread contagion to them more quickly.

For the simple network in **Figure 1**, Node B is the most central according to all three of these metrics. However, in other networks, different nodes might be most central according to different metrics. **Figure 2** depicts a more complicated network. This network features three nodes, or hubs, that have significantly more links than the other nodes. Node A has the most direct links among the three hubs (seven links compared with six for Node B and five for Node C), giving it the highest measure of degree centrality in the network. Because Node A has the most links, its impairment would directly disrupt the most interactions within the network. For instance, if **Figure 2** represented trading in an over-the-counter securities market, the impairment of institution A's operational systems would directly disrupt the most trades. As another example, a power outage at Atlanta's airport in December 2017 was particularly damaging because that's the world's busiest

Figure 1. Network Diagram and Adjacency Matrix

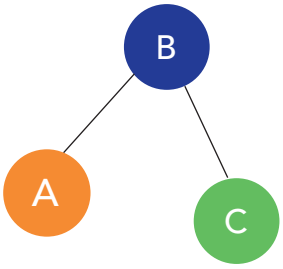

NODES	A	B	C
A		1	0
B	1		1
C	0	1	

Source: Authors' analysis

Figure 2. The Most Central Node by Different Metrics

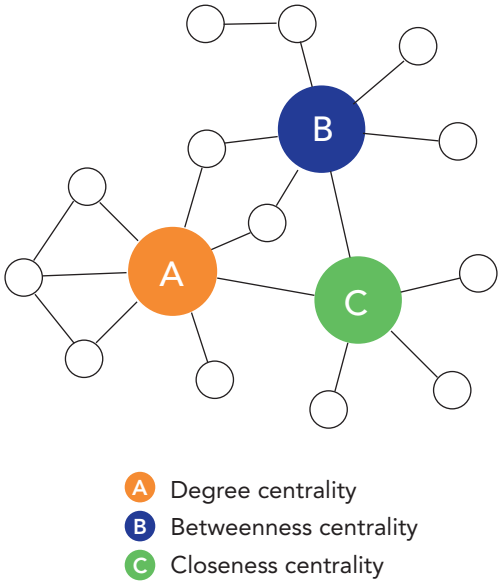

Source: Authors' analysis

airport. Hundreds of flights were cancelled because of the incident, affecting passengers in not only Atlanta, but also many other cities.[6] An incident at an airport with a lower degree centrality would not have affected as many flights or passengers, leading to less damage to the transportation network overall.

On the other hand, Node B is in more of the shortest paths that run between two nodes. Many of the nodes connected to Node B can access the rest of the network only through Node B, giving Node B the highest measure of betweenness centrality. This metric is important in networks in which information is transmitted across chains of nodes, such as a network for sharing credit ratings. An operational disruption to Node B could debilitate the network because rerouting the information flow through other paths could be costly or even impossible.

Node C, which is directly linked to Node A and Node B, can reach all other nodes in a small number of steps either through its own links or through these two hubs. Although Node C has the fewest direct links of the three hubs, it is the only hub directly linked to both other hubs, giving it the highest measure of closeness centrality. This metric is critical when considering cyber incidents that could spread from one firm's systems to another through contagion. In the presence of contagion, as in the 2003 electrical outage, this node would be critical to defend because it could spread contagion most quickly to other nodes in the network.

The Next Step: Defense Models

Although centrality metrics provide information about the criticality of different firms in a network, the questions they can address are limited. Network defense models help fill the gap. They best address questions about the damage cyber incidents can cause, how to best design financial networks, and the value of regulatory harmonization.

Network defense models build on basic network analysis by incorporating incidents against nodes, node defenses, and contagion. The standard model assumes a *designer* and an *adversary* are playing a game in which the adversary's effort can be targeted or random. This framework is general and can represent a variety of circumstances. For example, the designer could represent a regulator or a securities exchange owner, and the adversary could represent a malicious actor or an electrical power outage.

The designer can shape the network by setting up links, establishing defenses at nodes, or both. Better defenses at a node reduce the probability that an incident at the node would be damaging. Because the designer's resources are limited, fully defending all nodes would usually be too costly or technologically impossible. The designer must be strategic in choosing which nodes to defend.

After the designer sets up the defense structure, the adversary goes after nodes in the network, either at random or strategically. Affected nodes are assumed to be disabled. Through contagion, such nodes spread the impact to their neighbors. The overall damage done by the incident would depend on the network structure and defenses.

Contagion is a critical force to consider in establishing network defenses. The connections in the network can act as paths for threats. Many adversaries try to exploit these paths by infecting a few nodes and then going after other nodes linked to the initially infected ones. A seemingly minor initial incident can quickly become much more dangerous and destructive. For example, according to public reports, the NotPetya cyber incident in June 2017 began when an accounting software update was compromised.[7] The incident spread malware across the computer networks of organizations in more than 60 countries. The damage was enormous, and some affected companies each reported losses of more than $100 million.[8]

Different Threats, Different Defenses

The optimal defense structure varies depending on whether the anticipated incident is random or targeted. Random incidents hit nodes by chance. They may be acts of nature or the equivalent, such as hurricanes and electrical power outages, that cause operational failures and disable financial institutions.[9] Random incidents also could be the work of human adversaries that strike nodes at random because they have limited information about the most vulnerable node. In contrast, targeted incidents pinpoint the most vulnerable nodes

in a network. They might be the work of malicious actors who have detailed knowledge of the system and know which nodes to target to cause the most damage to certain firms or to the financial system.

The network designer must account for the nature of an anticipated incident when setting up network defenses. The designer can spend fewer resources on defense against random incidents by leaving some nodes relatively undefended because lower defenses do not increase the chance of such incidents. If the probability of an incident is low — for example, a node is in a region unlikely to be hit by a natural disaster — then leaving the node undefended could be more cost-effective.

However, networks with major vulnerabilities could be crippled by targeted incidents. Adversaries could exploit these vulnerabilities for maximum damage. Protecting every node against targeted incidents is typically not possible. Unless the designer has the resources to completely protect every node, the adversary could probably inflict damage by targeting less-protected nodes. The designer must structure the network defenses to minimize losses.

For example, suppose the adversary targeted a single node in the network in **Figure 2**. That node would fail unless protected. Through contagion, nodes directly linked with the first node would also fail unless protected. If the designer could protect only one node in the network from failure, which node would be the best to protect? The designer would aim to maximize the number of links remaining after the incident.

Figure 3 shows the average number of links that would remain if that node were the one protected from failure against a random incident that hit each node with equal probability. Node A would be the best to protect; Node B, the second best; and Node C, the third best. This ranking matches the ranking of the number of links each hub has initially (that is, the hub's degree centrality score).

In contrast, in a targeted incident the adversary may have detailed knowledge of the network's structure and defenses. In the example in **Figure 4**, the adversary is an enemy actor that always targets the node whose disruption would cause the most damage to the system. For instance, if Node A were protected,

Figure 3. Average Links Remaining After a Random Incident

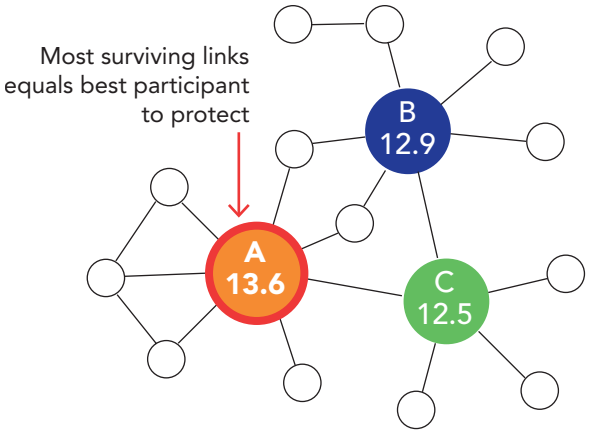

Source: Authors' analysis

Figure 4. Number of Links Remaining After a Targeted Incident

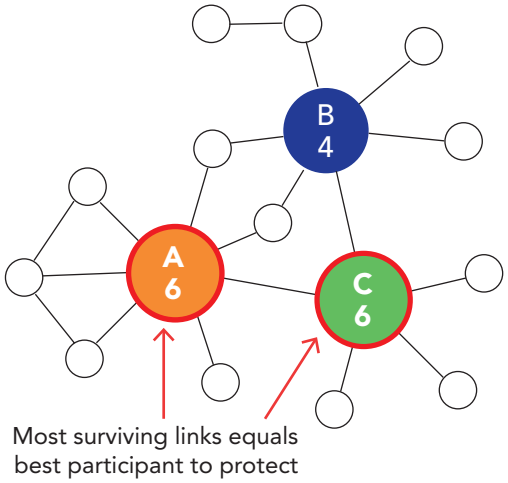

Source: Authors' analysis

targeting Node B would cause the most damage. Such a targeted incident would destroy the links of Node B and the links of all nodes directly connected to it, leaving six links in the network. On the other hand, if Node B were protected, targeting Node A would cause the most damage. The links of Node A and the links of all nodes directly connected to it would fail. Only four links would remain in the network. In either case, many fewer links would remain than after a random incident, an illustration of the additional risk targeted incidents can pose.

The relative ranking of hubs also differs for a targeted incident. Node C, which was the third best to protect from a random incident, is now tied with Node A for the most important node to protect from a targeted incident because it is connected to both other hubs. Contagion would give an adversary the opportunity to take out all three hubs by targeting Node C — the one with the highest level of closeness centrality. Closeness centrality is especially relevant if contagion is a risk.

This example shows the need to consider the type of shock when designing network defenses. In practice, random and targeted shocks are both possible. The relative likelihood of each type of shock would need to be weighed when choosing the appropriate defense structure.

Choosing the Network Structure

In the example in the previous section, the designer's job would be to erect defenses in an established network structure.[10] In some network defense modeling, the designer also determines how nodes are linked. That means the designer must weigh the strengths and weaknesses of possible structures.

One structure that can be particularly effective is the center-protected star network, or CP-star, named after its shape.[11] In a CP-star network, one node is at the center, and the other nodes are at the periphery. The central node is strongly protected, but peripheral nodes may be less protected. With the central node fully protected, contagion cannot spread when a peripheral node is impaired by an incident. The network will probably withstand an incident with limited losses. For example, in a financial setting, a CP-star structure might represent the market for standardized derivatives cleared through a single central counterparty. A node in this network would be an individual firm and its operational systems, and a link would be a submitted transaction. In such a market, the CCP's operational systems should be heavily fortified against cyber incidents. **Figure 5** shows a simple CP-star network. Node C is central, and the P nodes are peripheral.

A 2016 incident in which $81 million was stolen from Bangladesh's central bank shows both the strengths and weaknesses of a CP-star structure. In the incident, still-unidentified hackers penetrated Bangladesh Bank's computer system. They sent fraudulent payment messages that were authenticated through the international Society for Worldwide Interbank Financial Telecommunication (SWIFT) network. In this incident, only a peripheral node was compromised. The well-protected central node appears to have blocked further contagion. If the entire SWIFT system had been compromised, the result could have been much more damaging and perhaps systemic.

However, in SWIFT's CP-star structure, the peripheral nodes each have much of the responsibility for their own security, illustrating the potential downside of a decentralized approach to network defense. SWIFT did not impose strict security requirements on its members.

Figure 5. Example of a Center-protected Star Network

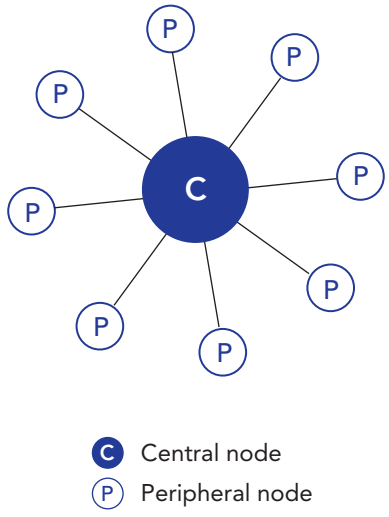

Source: Authors' analysis

Bangladesh Bank was following its own internal security protocols. Differences in security among institutions may explain why the core SWIFT system itself was not hacked, but Bangladesh Bank was. SWIFT has since implemented programs to improve its customers' security protocols.[12]

CP-star networks tend to be resilient, but they are not always the best choice. For example, if fully protecting the central node is not feasible, an incident impairing the center could debilitate the rest of the network. A designer trying to guarantee that some nodes survive such an incident might be better off separating the network into multiple unconnected parts.[13]

Figure 6 shows such a structure. Instead of a single central node to which all other nodes are connected, the network is broken into three distinct smaller networks, called *components.* An example of this structure would be a group of banks that each transact with different clients. With this defense structure, an incident that affects any single component cannot spread to others, so an incident is easier to contain. The disadvantage is that because not all nodes in the network are connected, not all nodes can transact with each other. This shortcoming may reduce functionality or increase costs. For instance, splitting up a CCP network among multiple CCPs could increase overall clearing costs because of less efficient netting of transactions.[14] When choosing the network structure, the designer must weigh such costs against gains in resilience.

Real-world networks are more complex than either a single CP-star or a multiple-components network of several CP-stars. However, such networks can still be analyzed using network defense models, and the lessons should still hold at a broad level.

Dangers of Decentralization

In many real-world settings, defense decisions are made by network participants independently from other participants, rather than by a single designer. For instance, financial institutions might set up their information technology systems according to their own internal procedures, consistent with any regulatory requirements. Lack of regulatory harmonization among regulators or countries can result in institutions

Figure 6. Example of a Multiple-components Network

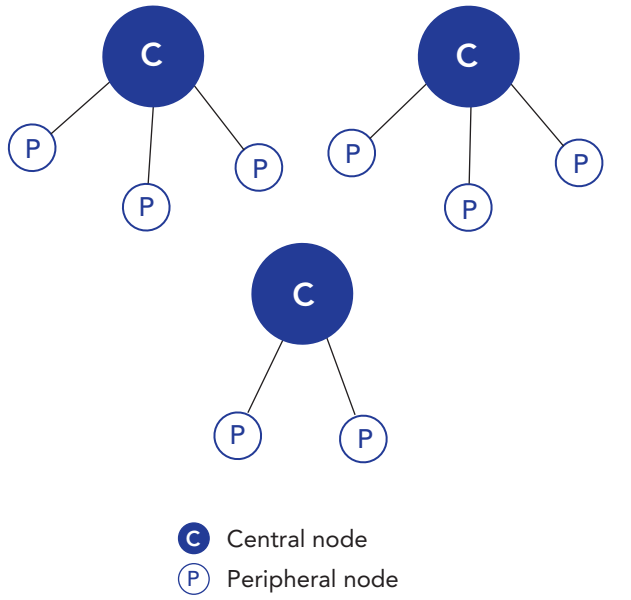

Source: Authors' analysis

having different security levels. If different firms are subject to different regulatory standards, the resulting system may suffer from a lack of coordination.

Decentralized decision making can lead to critical lapses in network defenses. Network defense models can be used to examine these issues. One way to incorporate decentralization into the models is to assume that a designer would choose which nodes are linked, but each node would choose its own defense investments.[15] For instance, a financial regulator could require central clearing of derivatives, which imposes a star network structure on the market. However, that regulator may not have the authority to set operational risk regulations for firms that trade in the network. Even though market structure and operational risk policies are closely related, they can be seen as separate issues within and among regulators.

A decentralized network defense strategy can result in underinvestment in defense if entities choose defense investments without regard for other participants. In the presence of contagion, an infected node will harm its neighbors, so investing in a strong defense benefits not only the node itself but also the node's neighbors. A firm that pays attention only to its own needs might

invest less on defense than would be beneficial for the system as a whole. In the aggregate, the result could be an insufficiently protected network.

Decentralization may also prevent the designer from using the CP-star structure even though it might be preferred if decisions were centralized.[16] As a stark example, suppose that the central node in **Figure 5** had the incentive to underinvest in its own security because it was not required to spend more. The entire network could then be vulnerable to collapse from an incident that hobbled this central node. In this case, the designer would need to separate the network into smaller pieces, as in **Figure 6**, to prevent incidents from propagating. In this way, a lack of coordination could have significant implications for optimal network design.

Conclusion

Large and interconnected networks are crucial to our modern financial system. The OFR uses network analysis to identify critical institutions in these networks and provide information that can help regulators. Network models produce useful insights for defending against operational risks. One insight is that different network structures may be required to defend against targeted versus random incidents. Another insight is that decentralized decisions about network defense can result in dangerous underinvestment in security.

Network models are one of several tools regulators have to address such issues. Regulators can combine network models with other methods and with empirical data as they perform testing. More detailed data and sophisticated models can allow for better results. Used appropriately, network models can assist in developing policies that enhance network resilience and support financial stability.

Awareness of network models also benefits industry participants. Firms should consider their roles within their networks, and analyze potential issues arising from contagion from neighbors or counterparties. With such awareness, firms could better set their security standards to protect against incidents. Greater cooperation and information sharing among firms would also mitigate the negative effects of decentralized decision making and improve system resilience.

Endnotes

[1] Stacey Schreft, Deputy Director for Research and Analysis (stacey.schreft@ofr.treasury.gov), and Simpson Zhang, Financial Economist, Office of the Comptroller of the Currency (simpson.zhang@occ.treasury.gov). This brief was written when Zhang was a Researcher at the OFR. The authors thank Brian Peretti, George Salmoiraghi, and Maryann Haggerty for their assistance.

[2] See J.R. Minkel, "The 2003 Northeast Blackout—Five Years Later," *Scientific American,* Aug. 13, 2008 (available at www.scientificamerican.com/article/2003-blackout-five-years-later, accessed April 25, 2018).

[3] See Stacy Cowley, "2.5 Million More People Potentially Exposed in Equifax Breach," *The New York Times,* Oct. 2, 2017 (available at www.nytimes.com/2017/10/02/business/equifax-breach.html?_r=0, accessed March 14, 2018).

[4] See, for example, Office of Financial Research, "Cybersecurity and Financial Stability: Risks and Resilience," OFR Viewpoint Paper no. 17-01, Feb. 15, 2017 (available at www.financialresearch.gov/viewpoint-papers/2017/02/15/cybersecurity-and-finanancial-stability, accessed March 14, 2018).

[5] See Mark Paddrik, Sriram Rajan, and H. Peyton Young, "Contagion in the CDS Markets," OFR Working Paper no. 16-12, Dec. 2, 2016 (available at www.financialresearch.gov/working-papers/2016/12/01/contagion-in-the-cds-market/, accessed April 30, 2018).

[6] See Russell Gold and Susan Carey, "Atlanta Airport Blackout Exposes a Flaw in Backup Power Systems," *The Wall Street Journal,* Dec. 19, 2017 (available at https://www.wsj.com/articles/atlanta-airport-blackout-exposes-a-flaw-in-backup-power-systems-1513638388, accessed April 25, 2018).

[7] See Jack Stubbs and Matthias Williams, "Ukraine Scrambles to Contain New Cyber Threat After 'NotPetya' Attack," Reuters, July 5, 2017 (available at www.reuters.com/article/us-cyber-attack-ukraine-backdoor/ukraine-scrambles-to-contain-new-cyber-threat-after-notpetya-attack-idUSKBN19Q14P, accessed March 14, 2018).

[8] See Scheherazade Daneshkhu, "Reckitt Seeks to Quantify Havoc of Malware Attack," *The Financial Times,* July 6, 2017 (available at www.ft.com/content/c4a63082-6264-11e7-91a7-502f7ee26895, accessed March 14, 2018).

[9] Random incidents do not require every node to be affected with the same probability. For example, some firms may be more at risk from hurricanes or earthquakes because of their geographic locations. The probability a node will be affected is fixed and not affected by the network's defense decisions.

[10] In some real-world settings, the designer may have no control over the links that agents form, and so such an assumption would be appropriate.

[11] This network was analyzed in a 2014 paper by Goyal and Vigier. See Sanjeev Goyal and Adrien Vigier, "Attack, Defence, and Contagion in Networks," *The Review of Economic Studies* 81, no. 4 (Oct. 1, 2014): 1518-1542 (available at https://academic.oup.com/restud/article-abstract/81/4/1518/1571797, accessed March 14, 2018).

[12] See Katy Burne and Robin Sidel, "Hackers Ran Through Holes in Swift's Network," *The Wall Street Journal,* April 30, 2017 (available at www.wsj.com/articles/hackers-ran-through-holes-in-swifts-network-1493575442, accessed Jan. 9, 2018); see also SWIFT and BAE Systems, *The Evolving Cyber Threat to the Banking Community.* Brussels: Society for Worldwide Interbank Financial Telecommunication, Nov. 29, 2017 (available at www.baesystems.com/en/cybersecurity/feature/the-evolving-cyber-threat-to-the-banking-community, accessed Jan. 3, 2018).

[13] The benefits of separating a network into multiple unconnected parts are shown in the model of Goyal and Vigier (2014).

[14] See for instance Darrell Duffie and Haoxiang Zhu, "Does a Central Clearing Counterparty Reduce Counterparty Risk?" *The Review of Asset Pricing Studies 1,* no. 1 (Dec. 1, 2011): 74-95 (available at https://academic.oup.com/raps/article/1/1/74/1528254, accessed March 14, 2018).

[15] As in the model of Cerdeiro, Dziubiński, and Goyal (2017). See Diego A. Cerdeiro, Marcin Dziubiński, and Sanjeev Goyal, "Individual Security, Contagion, and Network Design," *Journal of Economic Theory* 170 (July 2017): 182-226 (available at www.sciencedirect.com/science/article/pii/S0022053117300583?via%3Dihub, accessed March 14, 2018).

[16] A 2016 paper by Goyal and others modifies Goyal and Vigier (2014) and Cerdeiro and others (2017) by allowing players to choose both their links and defense choices. They show that this extra layer of decentralization can potentially lead to even more inefficient outcomes relative to Cerdeiro and others (2017). However, in simulations they show that the resulting networks actually tend to achieve reasonably high welfare. The network structures tend to feature multiple hub-spoke structures with connecting bridges. The authors also conduct a behavioral experiment with students, and they show that play in the experiment is quite similar to their theory and simulations. See Sanjeev Goyal and others, "Strategic Network Formation with Attack and Immunization," Dec. 11, 2016, at International Conference on Web and Internet Economics (available at https://link.springer.com/chapter/10.1007%2F978-3-662-54110-4_30, accessed March 14, 2018).