

Network Formation under Random Attack and Probabilistic Spread *

Yu Chen¹, Shahin Jabbari¹, Michael Kearns¹, Sanjeev Khanna¹ and Jamie Morgenstern²

¹University of Pennsylvania

²Georgia Tech

{chenyu2, jabbari, mkearns, sanjeev}@cis.upenn.edu, jamiemmt.cs@gatech.edu

Abstract

We study a network formation game where agents receive benefits by forming connections to other agents but also incur both direct and indirect costs from the formed connections. Specifically, once the agents have purchased their connections, an attack starts at a randomly chosen vertex in the network and spreads according to the independent cascade model with a fixed probability, destroying any infected agents. The utility or welfare of an agent in our game is defined to be the expected size of the agent’s connected component post-attack minus her expenditure in forming connections. Our goal is to understand the properties of the equilibrium networks formed in this game. Our first result concerns the edge density of equilibrium networks. A network connection increases both the likelihood of remaining connected to other agents after an attack as well the likelihood of getting infected by a cascading spread of infection. We show that the latter concern primarily prevails and any equilibrium network in our game contains only $O(n \log n)$ edges where n denotes the number of agents. On the other hand, there are equilibrium networks that contain $\Omega(n)$ edges showing that our edge density bound is tight up to a logarithmic factor. Our second result shows that the presence of attack and its spread through a cascade does not significantly lower social welfare as long as the network is not too dense. We show that any non-trivial equilibrium network with $O(n)$ edges has $\Theta(n^2)$ social welfare, asymptotically similar to the social welfare guarantee in the game without any attacks.

1 Introduction

We study a network formation game where strategic agents (vertices on a graph) receive both benefits and costs from forming connections to other agents. While various benefit functions exist in the literature [Bala and Goyal, 2000;

Fabrikant *et al.*, 2003], we focus on the *reachability network benefit*. Here, the benefit of an agent is the size of her connected component in the collectively formed graph. This models settings where reachability (rather than centrality) motivates joining the network, e.g. when transmitting packets over technological networks such as the Internet.

Most previous works feature a direct edge cost $c_e > 0$ for forming a link. Goyal *et al.* [2016] depart from this notion by studying a game where forming links introduces an additional *indirect* cost by exposing agents to contagious network shocks. These indirect costs can model scenarios such as virus spread through technological or biological networks.

Our work continues this investigation of direct and indirect connection costs. To model the indirect cost we assume that, after network formation, an adversary attacks a single vertex uniformly at random. The attack then kills the vertex and spreads through the network via the independent cascade model according to parameter p [Kempe *et al.*, 2003]. This random attack and probabilistic spread captures the epidemiological quality of virus spread in both biological and technological networks.

Our work is closely related to two previous works. Bala and Goyal [2000] study a reachability network game without attacks and show a sharp characterization of equilibrium networks: every tree and the empty network can form in equilibria. Goyal *et al.* [2016] study a reachability network formation game where an adversary inspects the formed network and then deliberately attacks a single vertex in the network. The attack then spreads deterministically to neighboring vertices according to a known rule, while agents may immunize against the attack for a fixed cost. Our game is most similar to the latter setting under a random adversary and high immunization cost. However, in our setting attacks spread probabilistically (through independent cascades) rather than deterministically. This yields an arguably more realistic model of infection spread but incurs additional complexity: computing the expected connectivity benefit of an agent in a given network is now #P-complete [Wang *et al.*, 2012].

Goyal *et al.* [2016] show that while more diverse equilibrium networks, including ones with multiple cycles, can emerge in addition to trees and the empty graph, the equilibrium networks with n agents will have at most $2n - 4$ edges; less than twice the number of edges that can form in the equilibria of the attack-free game. Furthermore, they

*The full technical version of this paper is available at <http://arxiv.org/abs/1906.00241>.

show that the social welfare is at least $n^2 - o(n^{5/3})$ in non-trivial equilibrium networks. Asymptotically, this is the maximum welfare possible which is achieved in any nonempty equilibrium of the attack-free game. In the regime where the cost of immunization is high, the game of Goyal et al. [2016] only admits disconnected and fragmented equilibrium networks due to deterministic spread of the attack, and the social welfare of the resulting networks may be as low as $\Theta(n)$.

Our Results and Techniques In our game, computing utilities or even verifying network equilibrium is computationally hard. We circumvent this difficulty by proving structural properties for equilibrium networks. First, we provide an upper bound on the edge density in equilibria.

Theorem 1 (Statement of Theorem 3). *Any equilibrium network on n vertices has $O(n \log n/p)$ edges.*

For constant p this upper bound is tight up to a logarithmic factor. The possibility of over-building therefore differentiates our game from those of Bala and Goyal [2000] and Goyal et al. [2016], but the extent of over-building is limited.

To prove Theorem 1, we first show that any equilibrium network with more than $\Omega(n \log(n/p))$ edges contains an induced subgraph with large minimum cut size. We then show that if a network has large minimum cut size, in every attack (with high probability), either almost all vertices in the network will die or almost all vertices in the network will survive. As a result, any vertex in the induced subgraph can beneficially deviate by dropping an edge. Together, these observations allows us to prove the claimed edge density bound.

Next, we show that any equilibrium network that is non-trivial (i.e. contains at least one edge) also contains a large connected component. Moreover, as long as the network is not too dense, it achieves a constant approximation to the best welfare possible of the attack-free game.

Theorem 2 (Informal Statement of Theorems 5 and 6). *Any non-trivial equilibrium network over n vertices contains a connected component of size at least $n/3$. Furthermore, if the number of edges in the network is $O(n/p)$, then the social welfare is $\Omega(n^2)$.*

To prove Theorem 2, we first show that any agent in a small connected component can increase her connectivity benefits by purchasing an edge to a larger component without significantly increasing her attack risk. This implies the existence of a large connected component. We then use the large component to argue that when the equilibrium network is sparse, the surviving network post-attack still contains a large connected component. This guarantees large social welfare.

While Goyal et al. [2016] show robustness of the structural properties of the original reachability game of Bala and Goyal [2000] to a variation with attack, deterministic spread and the option of immunization for players, we show robustness in another variant that involves a cascading attack but disallows immunization. However, on the technical front, the tools that we use to prove these robustness results are very different from the analysis of both of these previous games.

Organization We introduce our model and discuss the related work in Section 2. In Section 3 we present examples of equi-

librium networks of our game. Sections 4 and 5 are devoted to the characterization of the edge density and social welfare. We conclude with directions for future work in Section 6.

2 Model

We start by formalizing our model and borrow most of our notation and terminology from Goyal et al. [2016]. We assume the n vertices of a graph (network) correspond to individual players. Each player has the choice to purchase edges to other players at a fixed cost of $c_e > 0$ per edge. Throughout we assume that c_e is a constant independent of n . Furthermore, we use the term *high probability* to refer to probability at least $1 - o(1/n)$ henceforth.

A (pure) strategy $s_i \subseteq [n]$ for player i consists of a subset of players to whom player i purchased an edge. We assume that edge purchases are unilateral i.e. players do not need approval to purchase an edge to another player but that the connectivity benefits and risks are bilateral.¹

Let $\mathbf{s} = (s_1, \dots, s_n)$ denote the strategy profile for all the players. Fixing \mathbf{s} , the set of edges purchased by all the players induces an undirected graph. We denote a *game graph* as a graph G , where $G = (V, E)$ is the undirected graph induced by the edge purchases of all players.

Fixing G , the adversary selects a *single* vertex $v \in V$ uniformly at random to start the attack. The attack kills v and then spreads according to the independent cascade model with probability $p \in (0, 1)$ [Kempe et al., 2003].² In the independent cascade model, in the first round, the attack spreads independently killing each of the neighbors of the initially attacked vertex v with probability p . In the next round, the spread continues from all the neighbors of v that were killed in the previous round. The spread stops when no new vertex was killed in the last round or when all the vertices are killed.

The adversary’s attack can be alternatively described as follows. Fixing a game graph G , let $G[p]$ denote the *random* graph obtained by retaining each edge of G independently with probability p . The adversary picks a vertex v uniformly at random to start the attack. The attack kills v and all the vertices in the connected component of $G[p]$ that contains v .

Let $CC_i(v)$ denote the *expected* size of the connected component of player i post-attack to a vertex v and we define $CC_v(v)$ to be 0. Then the expected utility (utility for short) of player i in strategy profile \mathbf{s} denoted by $u_i(\mathbf{s})$ is precisely

$$u_i(\mathbf{s}) = \frac{1}{|V|} \sum_{v \in V} CC_i(v) - |s_i|c_e.$$

We refer to the sum of utilities of all the players playing a strategy profile \mathbf{s} as the *social welfare* of \mathbf{s} .

¹As an example of a scenario where the consequences are bilateral even though the link formation is unilateral, consider the spread of a disease in a social network where the links are formed as a result of physical proximity of individuals. The social benefits and potential risks of a contagious disease are bilateral in this case although the link formation as a result of proximity is unilateral. We leave the study of the bilateral edge formation for future work.

²Throughout we assume that p is a constant independent of the number of players n . We discuss the regime in which p decreases as the number of players increases in Section 4.1.

Wang et al. [2012] show that computing the exact spread of the attack in the independent cascade model is #P-complete in general. This implies that, given a strategy profile s , computing the expected size of the connected component of all vertices (and hence the expected utility of all vertices) is #P-complete. However, an approximation of these quantities can be obtained by Monte Carlo simulation.

We model each of the n players as strategic agents who deterministically choose which edges to purchase. A strategy profile s is a *pure strategy Nash equilibrium* if, for any player i , fixing the behavior of the other players to be s_{-i} , the expected utility for i , $u_i(s)$, cannot strictly increase when playing any strategy s'_i over s_i . We focus our attention to pure strategy Nash equilibrium (or equilibrium) in this work. Since computing the expected utilities in our game is #P-complete, even verifying that a strategy profile is an equilibrium is #P-complete. Hence as our main contribution, we prove structural properties for the equilibrium networks regardless of this computational barrier.

2.1 Related Work

There are two lines of work closely related to ours. First, Bala and Goyal [2000] study the attack-free version of our game. They show that equilibrium networks are either trees or the empty network. Also since there is no attack, the social welfare in nonempty equilibrium networks is asymptotically $n^2 - o(n^2)$.

Second, Goyal et al. [2016] study a network formation game where players in addition to having the option of purchasing edges can also purchase immunization from the attack. Since we do not study the effect of immunization purchases in our game, our game corresponds to the regime of parameters in their game where the cost of immunization is so high that no vertex would purchase immunization in equilibria. Moreover, they study several different adversarial attack models and our attack model coincides with their *random attack adversary*. The main difference between our work and theirs is that they assume the attack spreads deterministically while we assume the attack spreads according to the independent cascade model [Kempe *et al.*, 2003]. In many real world scenarios e.g. the spread of contagious disease over the network of people, the spread is *not deterministic*. Hence our work can be seen as a first attempt to make the model of Goyal et al. [2016] closer to real world applications. However, the change in the spread of attack comes with a significant increase in the complexity of the game as even computing the utilities of the players in our game is #P-complete. While Friedrich et al. [2017] have shown that best responses for players can be computed in polynomial time under various attack models, the question of whether best response dynamics converges to an equilibrium network is open in the model of Goyal et al. [2016].

Similar to Goyal et al. [2016] we show that diverse equilibrium networks can form in our game. While they show that all equilibrium networks over $n \geq 4$ players have at most $2n - 4$ edges, we show that the number of edges in any equilibrium network is at most $O(n \log n)$ and this bound is tight up to a logarithmic factor. Furthermore, Goyal et al. [2016] show

that the social welfare is asymptotically $n^2 - o(n^2)$ in non-trivial equilibrium networks. Their definition of non-trivial networks requires the network to have at least one immunized vertex and one edge. In the regime where the cost of immunization is high, the game of Goyal et al. [2016] only admits disconnected and fragmented equilibrium networks due to the deterministic spread of the attack. Such networks (even excluding the empty graph) can have social welfare as low as $\Theta(n)$. We show that any low density equilibrium network of our game enjoys a social welfare of $\Theta(n^2)$ as long as the network contains at least one edge.

Kliemann [2011] introduced a network formation game with reachability benefits and an attack on the formed network that destroys exactly one link with no spread. Their equilibrium networks are sparse and also admit high social welfare as removing an edge can create at most two connected components. Kliemann et al. [2017] extend this to allow attacks on vertices while focusing on swapstable equilibria.

Blume et al. [2011] introduce a network formation game with bilateral edge formation. They assume both edge and link failures can happen simultaneously but independent of the failures so far in the network. These differences make it hard to directly compare the two models. They show a tension between optimal and stable networks and exploring such properties in depth in our model is an interesting direction.

Finally, network formation games, with a variety of different connectivity benefit models, have been studied extensively in computer science see e.g. [Bala and Goyal, 2000; Blume *et al.*, 2011; Kliemann, 2011]. We refer the reader to the related work section of Goyal et al. [2016] for a comprehensive summary of other related work especially on the topic of optimal security choices for networks.

3 Examples of Equilibrium Networks

In this section we show that a diverse set of topologies can emerge in the equilibrium of our game. Similar to the models of Bala and Goyal [2000] and Goyal et al. [2016] the empty graph can form in the equilibrium of our game when $c_e \geq 1$. Moreover, similar to both models, trees can form in equilibria (See the left panel of Figure 1). Finally, while Goyal et al. [2016] show that in the regime of their game where the cost of immunization is high (so no vertex would immunize) no connected network can form in equilibria due to the deterministic spread of the attack, we show that connected networks indeed can form in the equilibria of our game (See Figure 1).




Figure 1: From left to right: hub-spoke, cycle and linear-paths network. A directed arrow determines the vertex that purchases the edge. We omit the details of the regime of parameters (c_e and p) in which such networks can form in the equilibria of our game.

We remark that pure strategy equilibria exist in all parameter regimes of our game. When $c_e \geq 1$, the empty network can form in equilibria for all p . When $c_e < 1$ a cycle or two disconnected hub-spoke structure of size $n/2$ can form in equilibria depending on whether p is far or close to 1 ($(1 - \omega(1/n))$ and $(1 - o(1))$, respectively).

Examples in Figure 1 show that denser networks can form in equilibria compared to the model of Bala and Goyal [2000] and the high immunization cost regime of the model of Goyal et al. [2016]. So it is natural to ask how dense equilibrium networks can be. We study this question in Section 4 and show an upper bound of $O(n \log n)$ on the density of the equilibrium networks. The examples in Figure 1 have $\Theta(n)$ edges, so our upper bound is tight up to a logarithmic factor.

Moreover, while all the equilibrium networks in Figure 1 are connected, there might still exist equilibrium networks in our game that are highly disconnected. In Section 5 we show that any equilibrium network with at least one edge contains a large connected component. However, even with the guarantee of a large connected component, there might still be concerns that the equilibrium networks can become highly fragmented after the attack. In Section 5 we show that as long as the equilibrium network is not too dense, the social welfare is lower bounded by $\Theta(n^2)$ i.e. a constant fraction of the social welfare achieved in the attack-free game.

We obtain these structural results even though we cannot compute utilities nor even verify that an equilibrium has reached due to computational barriers. We view these results as our most significant technical contributions.

4 Edge Density

We now analyze the edge density of equilibrium networks.

Theorem 3. *Any equilibrium network on n vertices has $O(n \log n/p)$ edges.*

The proof of Theorem 3 is due to the following observations which we formally state and prove next. At a high level, we first show that if G has large enough edge density, then G contains an induced subgraph H whose minimum cut size is large. We then show a large minimum cut size implies that $H[p]$ is connected with high probability. This means that in almost all attacks that infect a vertex in H , all vertices in H will get infected. So a vertex in H would have a beneficial deviation in the form of dropping an edge; which contradicts the assumption that G was an equilibrium network. This proves that equilibrium networks cannot be too dense.

More formally, we first show in Lemma 1 that if G is *dense enough* it contains a subgraph H with a minimum cut size, denoted by $\alpha(H)$, of at least $\Omega(\log n/p)$.

Lemma 1. *Let $G = (V, E)$ be a graph on n vertices. There exists a constant k such that if $|E| \geq kn \log n/p$ then G contains an induced subgraph H with $\alpha(H) \geq k \log n/p$.*

Proof. If $\alpha(G) \geq k \log n/p$ then G is the desired graph. Otherwise, there is a cut of size less than $k \log n/p$ that partitions G into two graphs G_1 and G_2 . Repeat this process at G_1 and G_2 , and build a decomposition tree T in this manner. Any leaf of this tree T is either a singleton vertex or a graph where the minimum cut size is at least $k \log n/p$. If at least one leaf

in T satisfies the latter property, then we are done and this is our desired graph H . We now argue that it can not be the case that all leaf vertices in T are singletons.

To see this, note that there can be at most $n - 1$ internal vertices in T and each internal vertex in T corresponds to removing up to $k \log n/p$ edges from G . Thus the decomposition process removes at most $k(n - 1) \log n/p$ edges. On the other hand, G has at least $kn \log n/p$ edges. It follows that not all leaves of T can be singleton vertices. \square

We then show that if $\alpha(G)$ is $\Omega(\log n/p)$ then with high probability $G[p]$ is connected.

Lemma 2 ([Alon, 1995]). *Let $G = (V, E)$ be a graph on n vertices. Then for any constant $b > 0$ there exists a constant $k(b)$ such that if $\alpha(G) \geq k(b) \log n/p$ then with probability at least $1 - n^{-b}$, $G[p]$ is connected.³*

We now define a property which we call *almost certain infection* and show that no equilibrium network can contain an induced subgraph satisfying this property.

Definition 1. *Let $G = (V, E)$ be a graph on n vertices and let H be a subgraph of G on more than one vertex. H has the almost certain infection property if whenever any vertex in H is attacked, then with probability at least $1 - o(1/n)$ the attack spreads to every vertex in H .*

Lemma 3. *Let $G = (V, E)$ be an equilibrium network on n vertices. G cannot contain an induced subgraph $H = (V', E')$ such that H satisfies the almost infection property.*

Proof. Consider any equilibrium graph G that violates the assertion of the claim. Let H be an induced subgraph of G with the almost certain infection property. We first prove that H contains a cycle. Assume by the way of contradiction that H does not have a cycle, so H is a collection of trees. Let u be any leaf in H . Then u is incident to at most one edge in H . Therefore, with probability $1 - p$, this edge is not in $G[p]$ and u is not connected to any other vertices in $G[p]$. So H cannot have the almost certain infection property (Recall that we assumed p is a constant independent of the number of players n). This means that H contains a cycle C . Let (u, v) be an edge on the cycle C . Assume without loss of generality that u purchased the edge (u, v) .

Now let ξ be the event that an attack propagates to some vertex in H after the attack. Then conditioned on ξ , with probability at least $1 - o(1/n)$, all vertices in H die. Hence vertex u in H has negative utility. On the other hand, if ξ does not occur, then the utility of u remains unchanged even if we remove the edge (u, v) . Thus vertex u can strictly improve her utility in this case by dropping the edge (u, v) . A contradiction to the fact that G is an equilibrium network. \square

We are now ready to prove Theorem 3.

Proof of Theorem 3. Assume by way of contradiction that G has more than $kn \log n/p$ edges where k is the constant in Lemma 1. Then by Lemma 1, G contains a subgraph $H = (V', E')$ such that $\alpha(H) \geq k \log n/p$. Since $|E'| \geq |V'|$,

³The statement in Alon [1995] requires $\alpha(G[p]) \geq k(b) \log n$. Since $\alpha(G[p]) = \alpha(G)p$ this translates to the condition stated in Lemma 2.

by Lemma 2, H has the almost certain infection property. However, G cannot be an equilibrium network by Lemma 3. \square

The most interesting regime for the probability of spread p is when p is a constant independent of n . While the upper bound in Theorem 3 holds for all p , it becomes vacuous as p gets small i.e. it becomes bigger than the trivial bound of $n^2/2$ when $p \leq k \log n/n$ for constant k . In Section 4.1 we analyze the edge density of equilibrium networks in the regime where $p < 1/n$. We show that the number of edges in any equilibrium network is bounded by $O(n)$ in this regime. To prove the density result we utilize properties of the Galton-Watson branching process and random graph model of Erdős-Rényi, as well as tools from extremal graph theory. We defer all the omitted details and proofs to the full version.

4.1 Small p Regime

In this section we focus on the regime where $p < 1/n$ and prove the following upper bound on the edge density.

Theorem 4. *Let $p = \kappa/n$ for some constant $\kappa < 1$. Let $G = (V, E)$ be an equilibrium network over n vertices. Then for sufficiently large n , $|E| \leq \max\{1/c_e, 24000\}n$.*

We prove Theorem 4 by contradiction and show that if the equilibrium graph has more than $\max\{1/c_e, 24000\}n$ edges, there exists a beneficial deviation in the form of dropping an edge for one of the players.

In order to prove Theorem 4, we need structural results stated in Lemmas 4 and 5. First, consider an edge (u, v) purchased by vertex u . Purchasing this edge would not have increased the connectivity benefit of u unless, after some attack, the edge (u, v) is the only path connecting u to v (and possibly other vertices that are only reachable through v). In Lemma 4 we show that if a graph is dense enough, then there exists an edge (u, v) such that many vertices should be deleted in order to make (u, v) the only remaining path connecting u and v .

Lemma 4. *Let $G = (V, E)$ be a graph on $n > 3\gamma$ vertices with $|E| \geq 2.5\gamma(n - \gamma) - 1$ for some $\gamma \in \mathbb{N}$. Then there exist two vertices v_1 and $v_2 \in V$ such that $(v_1, v_2) \in E$, and at least $\gamma + 1$ vertices need to be deleted so that the only path from v_1 to v_2 in G is through the direct edge (v_1, v_2) .*

Second, as described in Section 2, the number of vertices that are killed in any attack is the size of the connected component in $G[p]$ that contains the initially attacked vertex. Lemma 5 bounds the size of a randomly chosen connected component in $G[p]$.

Lemma 5. *Let $G = (V, E)$ be an equilibrium network over n vertices with $|E| = kn$ and $\max\{1/c_e, 24000\} \leq k = O(\log n)$. When $p < 1/n$ and n is sufficiently large, the size of the connected component of a randomly chosen vertex v in $G[p]$ is at most $k/3$ with probability at least $1 - 2c_e/(3n)$.*

We now give the formal proof of Theorem 4.

Proof of Theorem 4. Assume by way of contradiction that G is an equilibrium network with kn edges where $k > \max\{1/c_e, 24000\}$. Let $\gamma = \lfloor k/2.5 \rfloor$, we have $kn > 2.5\gamma(n - \gamma) - 1$. By Lemma 4, there exists an edge (u, v) such

that in order to make (u, v) the only remaining path connecting u and v , we need to delete at least $\gamma + 1$ vertices. Without loss of generality assume that u has purchased the edge (u, v) . If in any attack at most γ vertices are killed, u will not lose any connectivity benefit after dropping this edge but decrease her expenditure by c_e .

Consider the size of the largest connected component in $G[p]$. Since $\kappa < 1$, the size of the largest connected component in $G[p]$ is at most $\beta \log n$ with probability $1 - o(1/n)$ for sufficiently large constant β which only depends on κ . When G is the complete graph, $G[p]$ corresponds to the random graph generated by the Erdős-Rényi model. In such case, the size of the largest component of $G[p]$ is $O(\log n)$, with high probability, when $pn = \kappa < 1$ [Erdős and Rényi, 1960].

If $\gamma > \beta \log n$, then with probability at most $o(1/n)$, the attack kills more than γ vertices, in which case the connectivity benefit of u can decrease by at most n after dropping the edge (u, v) . So the expected connectivity benefit of u decreases by at most $o(1)$ after the deviation but her expenditure also decreases by c_e . Hence, after the deviation, the expected change in the utility of u is at least $c_e - o(1) > 0$ which contradicts the assumption that G is an equilibrium network (recall that we have assumed c_e is a constant independent of n).

If $\gamma \leq \beta \log n$, then by definition, $k = O(\log n)$. Since k is also at least $\max\{24000, 1/c_e\}$, by Lemma 5, with probability at least $1 - 2c_e/(3n)$, the attack kills at most $k/3 < \gamma$ vertices, in which case the connectivity benefit of u remains unchanged after dropping the edge (u, v) . With probability at most $2c_e/(3n)$, more than $k/3$ vertices are killed in which case the connectivity benefit of u can decrease by at most n . So the expected connectivity benefit of u decreases by at most $2c_e/3$ after the deviation but her expenditure also decreases by c_e . Hence, after the deviation, the expected change in the utility of u is at least $c_e/3 > 0$, which contradicts the assumption that G is an equilibrium network. \square

5 Social Welfare

In this section we provide a lower bound on the social welfare of equilibrium networks. Similar to other reachability games, the empty graph can form in equilibrium [Bala and Goyal, 2000; Goyal *et al.*, 2016]. Hence without any further assumptions, no meaningful guarantee on the social welfare can be made. Hence, we focus on non-trivial equilibrium networks defined as follows.

Definition 2. *An equilibrium network is non-trivial if it contains at least one edge.*

Definition 2 rules out the empty network but it is still possible that a non-trivial equilibrium network contains many small connected components or becomes highly fragmented after the attack. In this section we show that none of these concerns materialize. In particular, in Theorem 5, we first show that any non-trivial equilibrium network contains at least one large connected component. We then show in Lemma 6 that when the network is not too dense, the equilibrium network cannot become highly fragmented after the attack. These two observations allows us to prove our social welfare lower bound as stated in Theorem 6.

We start by showing that any non-trivial equilibrium network contains a large connected component.

Theorem 5. *Let $G = (V, E)$ be a non-trivial equilibrium network over n vertices. Then, for sufficiently large n , the largest connected component of G has at least $n/3$ vertices.*

We next present Lemma 6 that describes the relationship between the expected size of the largest connected component of $G[p]$ and the connectivity benefits of the vertices in G .

Lemma 6. *Let $G = (V, E)$ be an equilibrium network over n vertices. Let C be any connected component in G of size n_C . If the expected size of the largest component of $C[p]$ is at most $n_C(1 - \epsilon)$, then the expected sum of the connectivity benefits of the vertices in C is at least $(n - n_C)n_C^2/n + \epsilon n_C^3/(3n)$.*

Theorem 5 and Lemma 6 allow us to prove a lower bound on the social welfare of non-trivial equilibrium networks.

Theorem 6. *Let $G = (V, E)$ be a non-trivial equilibrium network over n vertices. For any $\epsilon \in (0, 1/8)$ and sufficiently large n , if $|E| < (1 - 2\epsilon)n \log(1/\epsilon)/p$, then the social welfare of G is at least $\epsilon n^2/3 - O(n/p)$.*

Proof. We show the expected sum of the connectivity benefits of the vertices in G is at least $\epsilon n^2/3$. Subtracting off the cumulative expenditure for edge purchases which is $|E|c_e = O(n/p)$ then imply the statement of the theorem.

Suppose the largest connected component of G say C^* has size n^* . By Theorem 5, $n^* \geq n/3$. We consider two cases based on the size of n^* : (1) $n^* \leq (1 - 2\epsilon)n$ and (2) $n^* > (1 - 2\epsilon)n$.

In case (1), where $n^* \leq (1 - 2\epsilon)n$, the sum of the connectivity benefits of the vertices in the largest connected component is at least $(n - n^*)n^{*2}/n$. This is because with probability of $(n - n^*)/n$ the attack starts outside of the component and all the vertices in the component survive; in such case the sum of connectivity benefits of the vertices in C^* is n^{*2} . Moreover, the derivative of $(n - n^*)n^{*2}/n$ with respect to n^* is $(-3n^{*2} + 2nn^*)/n$, which is positive when $0 < n^* < 2n/3$ and negative when $n^* > 2n/3$. Since $n^* \in [n/3, (1 - 2\epsilon)n]$, the minimum value of $(n - n^*)n^{*2}/n$ should be at one of the end points which correspond to values $2n^2/27$ or $2\epsilon(1 - 2\epsilon)^2n^2$, respectively. Both of these values are larger than $\epsilon n^3/3$ when $\epsilon < 1/8$, which means the sum of connectivity benefits is at least $\epsilon n^2/3$ in this case.

In case (2), where $n^* > (1 - 2\epsilon)n$, the number of edges in the connected component C^* is at most $n^* \log(1/\epsilon)/p$ (which occurs when all the edges are in this component). Let us denote the vertices in C^* by numbers from 1 to n^* . Let d_i denote the degree of vertex i . We first bound the expected number of isolated vertices of C^* in $G[p]$. A vertex becomes isolated in $G[p]$ if none of the edges adjacent to it are sampled to be retained. This event occurs with probability $(1 - p)^{d_i}$ for vertex i . So we can derive a lower bound on the expected number of isolated vertices of C^* in $G[p]$ as follows.

$$\sum_{i=1}^{n^*} (1 - p)^{d_i} \geq n^* (1 - p)^{\frac{2|E|}{n^*}} > \epsilon n^*,$$

where the first inequality is by inequality of arithmetic and geometric means and the second inequality is by the assumption that $|E| < n^* \log(1/\epsilon)/p$. Thus the expected size of the

largest connected component in $C^*[p]$ is at most $n^*(1 - \epsilon)$. We can now apply Lemma 6 to show that the expected sum of the connectivity benefits of the vertices in C^* is at least

$$\frac{\epsilon n^{*3}}{3n} + \frac{(n - n^*)n^{*2}}{n} = \frac{n^{*2}n - (1 - \epsilon/3)n^{*3}}{n}.$$

This is strictly decreasing in n^* as the derivative with respect to n^* is negative. So the expected sum of connectivity benefits of the vertices in C^* (and hence in G) is at least $\epsilon n^2/3$ (when $n^* = n$). \square

Finally, we remark that unlike the models of Bala and Goyal [2000] and Goyal et al. [2016], achieving a social welfare of $n^2 - o(n^2)$ is impossible in our game even when restricting to sparse and non-trivial equilibrium networks. This is formalized in Proposition 1.

Proposition 1. *There exists a non-trivial equilibrium network $G = (V, E)$ over n vertices with $O(n)$ edges such that the social welfare of G is kn^2 for $k < 1$.*

Proof. The hub-spoke equilibrium (see Figure 1) with $p = 0.6$ satisfies the condition of Proposition 1 with $k = 0.4$. \square

6 Conclusions

We studied a natural network formation game where each network connection has the potential to both bring additional utility to an agent as well add to her risk of being infected by a cascading infection attack. We showed that the equilibria resulting from these competing concerns are essentially sparse and containing at most $O(n \log n)$ edges. We also showed that any non-trivial equilibrium network in our game achieves the highest possible social welfare of $\Theta(n^2)$ whenever the equilibrium network has only $O(n)$ edges.

The Price of Anarchy in our model is $\Theta(n)$. To illustrate, consider the $c_e \geq 1$ regime. A central planner can build a cycle or two disconnected hub-spoke structures of size $n/2$ depending on whether the probability of spread of the attack p is low or high, respectively (and both of these structures can also form in equilibrium). Such networks have social welfare of $\Theta(n^2)$. However, the empty network is an equilibrium network in this regime implying a Price of Anarchy of at least $\Theta(n)$ – the worst Price of Anarchy possible. The Price of Stability in our model is $\Theta(1)$ since the social welfare is trivially bounded by n^2 and either of the two equilibrium networks above achieve a social welfare of $\Theta(n^2)$.

Our results suggest several natural questions for future work. Our upper bound of $O(n \log n)$ is a logarithmic factor higher than the densest equilibrium network that we can create. Narrowing this gap is the most interesting open question. Improving our network density upper bound to $O(n)$ edges would immediately imply that all non-trivial equilibrium networks achieve $\Omega(n^2)$ social welfare. Another direction for future work is to analyze how network density and social welfare evolves when agents additionally have an option to invest in immunization that protects them from infections.

References

- [Alon, 1995] Noga Alon. A note on network reliability. *Discrete Probability and Algorithms*, pages 11–14, 1995.
- [Bala and Goyal, 2000] Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1230, 2000.
- [Blume *et al.*, 2011] Larry Blume, David Easley, Jon Kleinberg, Robert Kleinberg, and Éva Tardos. Network formation in the presence of contagious risk. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 1–10, 2011.
- [Erdős and Rényi, 1960] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, 5(1):17–60, 1960.
- [Fabrikant *et al.*, 2003] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos Papadimitriou, and Scott Shenker. On a network creation game. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing*, pages 347–351, 2003.
- [Friedrich *et al.*, 2017] Tobias Friedrich, Sven Ihde, Christoph Keßler, Pascal Lenzner, Stefan Neubert, and David Schumann. Efficient best response computation for strategic network formation under attack. In *Proceedings of the 10th International Symposium on Algorithmic Game Theory*, pages 199–211, 2017.
- [Goyal *et al.*, 2016] Sanjeev Goyal, Shahin Jabbari, Michael Kearns, Sanjeev Khanna, and Jamie Morgenstern. Strategic network formation with attack and immunization. In *Proceedings of 12th International Conference on Web and Internet Economics*, pages 429–443, 2016.
- [Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.
- [Kliemann *et al.*, 2017] Lasse Kliemann, Elmira Shirazi Sheykhdarabadi, and Anand Srivastav. Swap equilibria under link and vertex destruction. *Games*, 8(1):14, 2017.
- [Kliemann, 2011] Lasse Kliemann. The price of anarchy for network formation in an adversary model. *Games*, 2(3):302–332, 2011.
- [Wang *et al.*, 2012] Chi Wang, Wei Chen, and Yajun Wang. Scalable influence maximization for independent cascade model in large-scale social networks. *Data Mining and Knowledge Discovery*, 25(3):545–576, 2012.