

Network Formation in the Presence of Contagious Risk ^{*}

Larry Blume [†] David Easley [‡] Jon Kleinberg [§]
Robert Kleinberg [¶] Éva Tardos ^{||}

Abstract

There are a number of domains where agents must collectively form a network in the face of the following trade-off: each agent receives benefits from the direct links it forms to others, but these links expose it to the risk of being hit by a cascading failure that might spread over multi-step paths. Financial contagion, epidemic disease, and the exposure of covert organizations to discovery are all settings in which such issues have been articulated.

Here we formulate the problem in terms of strategic network formation, and provide asymptotically tight bounds on the welfare of both optimal and stable networks. We find that socially optimal networks are, in a precise sense, situated just beyond a phase transition in the behavior of the cascading failures, and that stable graphs lie slightly further beyond this phase transition, at a point where most of the available welfare has been lost. Our analysis enables us to explore such issues as the trade-offs between clustered and anonymous market structures, and it exposes a fundamental sense in which very small amounts of “over-linking” in networks with contagious risk can have strong consequences for the welfare of the participants.

^{*}Supported in part by the MacArthur Foundation, a Google Research Grant, a Yahoo! Research Alliance Grant, NSF grants CCF-0325453, IIS-0910664, CCF-0910940, and CCF-0729006, and ONR grant N00014-98-1-0589.

[†]Dept. of Economics, Cornell University, Ithaca NY 14853; IHS, Vienna; and the Santa Fe Institute. Supported in part by WWTF Grant Die Evolution von Normen and Konventionen in der Wirtschaft.

[‡]Dept. of Economics, Cornell University, Ithaca NY 14853.

[§]Dept. of Computer Science, Cornell University, Ithaca NY 14853.

[¶]Dept. of Computer Science, Cornell University, Ithaca NY 14853.

^{||}Dept. of Computer Science, Cornell University, Ithaca NY 14853.

1 Introduction

There are many situations in which self-interested agents form links with one another, producing an underlying network structure. A growing body of work on *strategic network formation* [21] has sought to analyze the structure of networks that develop through such types of interaction: what characteristic properties do they have, and are they approximately as efficient as networks created through central planning and coordination? A common theme running through this literature is the view that links are costly to form, and this cost must be traded off against the benefits of indirect access to other parts of the network through multi-step paths, as measured by distances [12, 13, 18], component sizes [7], or point-to-point connectivity [6].

In a wide range of settings where network formation occurs, however, this trade-off is inverted. Instead of costly links that confer indirect benefits, the agents receive benefits via their direct links, but through these links they risk exposure to negative payoffs from a broader set of nodes reachable on multi-step paths. A canonical example is the problem of financial contagion [2, 3, 11, 16]. In a highly simplified formulation of this problem, two parties A and B engaging in a transaction must balance the benefits arising from the intended successful outcome of their transaction against the underlying *counterparty risk*: the possibility that the other side will fail to complete the transaction. This risk can now spread across the links formed by multiple transactions: if A defaults on its agreement with B , this may mean that B lacks the resources needed to complete a concurrent transaction with C , and so B too defaults. In this way the failure of A has spread via B to affect a transaction involving C , and perhaps further. The fear of such cascading failures, and the behavior of financial institutions in response to this perceived increase in counterparty risk, played an important role in the financial crisis of Fall 2008.

Financial contagion, however, is just one setting in which strategic agents balance the benefits of linking against concerns about contagious risk. Other basic examples include the following.

- During a disease epidemic within a human population, high-risk groups will alter their interaction patterns to reduce the risk of contagion, including focusing more on “in-group” members who belong to similar social circles. In the setting of HIV/AIDS, for example, this has been studied in the context of sexual contacts [19] and needle-sharing among intravenous drug users [8].

- In any type of distributed file-sharing application, the spread of malware from one host to another creates a very similar form of contagious risk, which can be incorporated into decisions about which hosts to download content from.
- In a different setting, participants in a clandestine organization benefit from the links they form to other members for purposes of coordination and mutual assistance, but by forming such links they expose themselves to a form of contagious risk: if certain members are compromised by an adversary, they may in turn reveal the identities and whereabouts of others to whom they are linked, and this process could potentially spread for multiple steps through the organizational network [15].

Formulating a Model

In all these different settings, the precise formulation of the payoffs to the participants will depend on details of the domain — for example, precisely how does a node benefit from the presence of a link, and precisely how does the risk of negative payoff propagate from one node to another? One of the advantages that general models of strategic network formation have offered is the ability to draw qualitative insights about the effects of self-interested behavior on network structure across a range of scenarios, using stylized formulations that capture the basic trade-offs without delving into the idiosyncracies of any one domain.

A natural modeling strategy for studying strategic network formation is to define a non-cooperative game whose outcomes are graphs. The precise networks that will emerge will depend on the precise details of the game. An enumeration of the possibilities is huge, and the likelihood of finding tractable formulations seems low. Network theorists, following [18], have finessed this problem by identifying properties of networks that we could expect to be satisfied by all equilibrium networks of games of conceivable interest.

In this paper, we develop a general model in this spirit to capture the underlying trade-off between the benefit of link formation and the problem of contagious risk, using simple definitions for the payoffs arising from these underlying processes. The model is formulated as follows. To begin with, we have a set V of n players, and players can choose to form bilateral relationships with one another, resulting in an undirected graph $G = (V, E)$. A player receives a payoff of $a > 0$ from each relationship it takes part in. We will generally assume there is an upper bound Δ on the number of links any one node is able to form.

Now, once the network is formed, a random process creates cascading failures as follows. First nodes fail independently with probability q , and then failed nodes have a probability of p of causing their neighbors to fail as well, with the failure potentially continuing to spread from these newly failed nodes. In more detail:

- First, each player randomly experiences a failure, independently with probability $q > 0$. We will refer to these as the *root failures* in the graph.
- Next, we declare each edge of G to be *live* independently with probability p and *blocked* with probability $1 - p$. We think of the live edges as those that transmit failure, and the blocked edges as those that do not transmit failure. Any node that can reach a root failure using a path consisting entirely of live edges is declared to fail also.

If a player fails, it loses any benefit from the links it forms, and instead its payoff is defined to be $(-b) < 0$.

Thus, if d_i denotes the degree of node i in G , and ϕ_i denote the probability that it fails (taken over the random choices of root failures and live edges), then we can write i 's expected payoff as

$$\pi_i = ad_i(1 - \phi_i) - b\phi_i = ad_i - (ad_i + b)\phi_i.$$

As noted above, our formulation of the payoffs is intended to capture the basic trade-off in a simple way. Links confer benefits that scale linearly in the degree, and failures spread through direct probabilistic contagion across edges. One can imagine more complex models for both of these aspects of the payoff, with more complex notions of the way in which a node's links increase its payoff, and more complex mechanisms for the spread of failures. These are interesting directions for further analysis. Here we will see that the present model already exhibits rich behavior, and suggests avenues for pursuing such generalizations.

Optimality and Stability. For most of the discussion, we will measure the “quality” of a graph via its *minimum welfare* (henceforth abbreviated *min-welfare*), which is the minimum payoff of any node in the graph. Unless otherwise specified, a socially optimal graph is one that maximizes this quantity.

In addition to socially optimal graphs, we will also consider graphs that arise under a solution concept representing strategic behavior by the nodes. We say that a graph is *stable* if (i) no node can strictly increase its payoff by deleting all its incident links (hence removing itself from the network), and (ii) there is no pair

of nodes (i, j) such that (i, j) is not an edge of G , but both i and j would have higher payoffs, with at least one of them strictly higher, if (i, j) were added to G . This notion of stability is in the spirit of solution concepts for the formation of undirected graphs, capturing the notion that it takes two nodes to agree on the formation of a link, but any node can unilaterally withdraw from its links. Our definition of stability is a relaxation of the notion of *pairwise Nash stability* [17], which modifies (i) to allow a node to drop any subset of its incident links. Thus, any pairwise Nash stable graph is also stable under our definition, and so when we show upper bounds on the welfare of all stable graphs, it applies to all pairwise Nash stable graphs as well.

When we consider the structures of socially optimal and stable graphs, much of the interesting behavior emerges in a natural range of the parameters a , b , p , and q motivated by the following considerations. Suppose we had just two nodes i and j , and suppose that i is deciding whether to link to j . If it forms the link, it receives a benefit of a but there is a probability of pq that j will fail and that this failure will spread to a . We want i to be willing to form the link to j under these conditions, and so we assume $a > bqp$. Otherwise no links will form. On the other hand, suppose that i knew that j were going to fail, so that the only thing protecting i from failure is the transmission probability p . Under these conditions we do not want i to form the link to j , so we assume $a < bp$. Otherwise there will be no strategic component to the analysis. Analogously, suppose that i knew that any failure at j would definitely spread to i , so that the only thing protecting i from failure is the chance $1 - q$ that j does not fail. Under these conditions we also do not want i to form the link to j , so we assume $a < bq$.

In our analysis, we will focus on the range of parameters in which these bounds hold by arbitrarily large constant factors. That is, we will consider the case in which p and q are small, and for some sufficiently large constant d , the quantity a exceeds bqp by at least a factor of d , and $\min(bp, bq)$ in turn exceeds a by at least a factor of d . So with $\delta = 1/d$, we have $\delta^{-1}bqp < a < \delta \min(bp, bq)$; we will refer to δ as the key *separation parameter* in our model. Finally, we will consider the case in which the number of nodes n is arbitrarily large compared to these other quantities (and/or their reciprocals).

Our Results

Our main results provide asymptotically tight characterizations of the welfare obtained by both socially optimal and stable graphs, as well as structural insights into the properties of such graphs. In a sense to be made precise below, we find

roughly that social optimality occurs just beyond the edge of a phase transition that controls how failures propagate, while stable graphs lie slightly further still past this phase transition, at a point where most of the welfare has already been wiped out. This narrow region where optimality occurs, and the inability of strategic agents to stay within this region, is one of the key qualitative behaviors to emerge from our model.

Social Optimality. For social optimality, our first main result is that there is a small $\varepsilon > 0$, going to 0 with our separation parameter δ , such that no graph can have min-welfare greater than $(1 + \varepsilon)a/p$. We show this by establishing that once each node forms more than $1/p$ links, the live edges — those that transmit failure — form a large connected component with high probability; this causes the probability of node failure to rise abruptly and to correspondingly reduce the welfare. This phenomenon is an analogue of the giant-component phase transition in the traditional random graph model $\mathcal{G}(n, p)$ [10], but since we are dealing with the random choice of live edges in an arbitrary underlying graph defined by the agents, we need to study the random subgraphs of this arbitrary graph; and hence we must adapt the arguments about the emergence of giant components to this setting.

This analysis establishes node degrees of $1/p$ as an important critical point in the behavior of the nodes' payoffs: at a degree of $1/p$, a node is receiving a benefit of a/p , and in order to raise all the nodes' payoffs further, we must push the graph into a region where the propagation of failures via live edges increases dramatically. It is therefore useful to look more carefully at the structure of graphs and their resulting payoffs right at the edge of this region. Accordingly, we say that a family of graphs, with n going to infinity, has *super-critical payoffs* if their min-welfares exceed a/p by a multiplicative constant (necessarily depending on the separation parameter δ) that is strictly greater than 1 — in other words, if they have min-welfare at least $(1 + f(\delta))a/p$. Otherwise, we will say that they have *sub-critical payoffs*. It then becomes an interesting question to consider whether there exist families of graphs achieving super-critical payoffs, and what can be said about their structure.

Clustered vs. Anonymous Markets. We construct a family of graphs with super-critical payoffs, by having the agents form a union of disjoint cliques, where each clique has a size chosen to be very slightly above $1/p$. The analysis of such graphs in fact highlights an interesting qualitative contrast between two kinds of

network structures in which the agents can be organized: (i) a union of disjoint cliques each of size $k + 1$, or (ii) a random graph formed by linking each node to k other nodes selected uniformly at random. Viewing this contrast in terms of an underlying motivation via economic contagion, we see that the two network structures correspond to two very different kinds of market formations: (i) is a *clustered market* where people engage in mutual dealing within their own community, and thereby insulate themselves from trouble originating in other communities, while (ii) is an *anonymous market* where agents can only specify the *number* of counterparties they would like to deal with, but cannot specify *who* these counterparties will be.

What our analysis shows is that while it is possible for clustered markets to yield super-critical payoffs, anonymous market structures can only produce sub-critical payoffs. In other words, if we connect each agent to k other random agents, then once k reaches $(1 + \gamma)/p$ for any positive constant γ (and n sufficiently large), the node payoffs become sub-optimal. This yields potential insights into the different contagion dynamics of the two kinds of structures — when the market institution enables people to separate themselves into protected communities, then it becomes feasible to form a number of links that pushes past the crucial phase transition; but when interactions happen anonymously, it is necessary to stay safely on the low side of the phase transition.

Stable Graphs. We then consider the class of stable graphs, and here our main result is that there is a small $\varepsilon > 0$ (again going to 0 with the separation parameter δ) such that no stable graph can have min-welfare greater than $\varepsilon a/p$. It is interesting that this gap between welfare at social optimality and at stability arises because of relatively subtle structural differences between the two kinds of graphs; socially optimal graphs are poised at the edge of a dangerous precipice in parameter space, and intuitively, we can think of stable graphs as exhibiting a small amount of additional linking that send them over this precipice.

Our analysis provides a useful way of thinking about how this additional linking occurs. When a node evaluates forming an extra link, its own change in payoff comes from the benefit of this new link offset by the increased risk of failure. However, while only the two endpoints of the new link can experience the benefits of the link, a potentially much larger set of nodes can experience a payoff decrease due to greater risk. This creates a negative externality by which nodes do not account for the full amount of risk they generate in forming links, and since we are in a region of parameter space where small structural changes can yield

large payoff changes, this externality is sufficient to eliminate most of the payoff to nodes.

Finally, as part of our analysis, we also provide a proof that stable graphs exist, and in fact that one can build stable graphs from unions of disjoint cliques. This construction requires some amount of care, since the stability condition requires that we produce cliques large enough that nodes will not want to form links into other cliques, but not so large that nodes will want to drop out of their own cliques. Arguing that there exists a clique size achieving both of these conditions involves maintaining control of the payoff functions in the vicinity of the phase transition.

Further Related Work

There have been several papers that deal with problems where one optimizes parameters of a network under the risk of node failure, but these other models have been quite different from ours both in the general questions they consider and in the specifics of their formulations.

Caballero and Simsek consider a model of financial contagion in which the failure of nodes in a network can lead other nodes to engage in a cascade of “fire sales,” in which the price of assets plummets [11]. However, in addition to the fact that their cascade mechanism is quite different, they study only the case in which the network structure is simple and given: the nodes are assumed to be connected in a cycle. Goyal and Vigier consider “attack-defense” games on networks, in which a designer must choose a network structure, and then an attacker tries to disrupt it by destroying nodes [14]. Their work assumes that failure spreads deterministically between nodes, and hence leads to very different considerations. Gutfraind [15], motivated by the organization of terrorist cells, considers the question of how to centrally design a network that can optimize a distance-based function in the presence of probabilistic contagion; in his work, the objective function is again very different from ours, and the design of such networks is explored through simulation. Finally, the line of research beginning with Albert, Jeong, and Barabási on the attack-tolerance of random graphs [1] focuses on the increase in average node distances when an adversary deletes nodes from an underlying graph, and it leads to different types of questions as well.

2 An Upper Bound on the Optimal Min-Welfare

We begin by establishing an upper bound on the min-welfare of any graph. We use the parametrization of the model described in the introduction; recall that we will be making the following assumption (motivated there): for a small constant $\delta > 0$, we have

$$\delta^{-1}bqp < a < \delta \min(bp, bq).$$

For ease of future reference, we will call this *Assumption* $\mathcal{P}(\delta)$. Recall also that we are assuming an upper bound of Δ on the maximum number of edges any one node can form. As discussed above, much of the interesting behavior in this model turns out to take place in graphs where the average degree is close to $1/p$. As a result, we will assume that Δ is larger than $1/p$, but not so large that any single node can dominate the structure of the graph. In particular, we will assume that $\Delta = c^*/p$ for a constant $c^* > 1$.

Our basic plan for the upper bound is as follows. If the min-welfare in a graph G exceeds $\frac{(1+\varepsilon)a}{p}$, then in particular it implies that all node degrees are at least $\frac{(1+\varepsilon)}{p}$. We declare each edge to be *live* independently with probability $p > 0$, and then study the extent to which nodes are able to reach one another along *live-edge paths* — that is, paths consisting entirely of live edges. If there is a node that can reach many others with reasonable probability, this node experiences a large probability of failure, and hence has a sharply reduced payoff, which will ultimately contradict our assumption that G has large min-welfare.

Now, how do we show that some node has a reasonably high chance of reaching many others on live-edge paths? There is a connection to the basic random graph model $\mathcal{G}(n, p)$, in which an edge is inserted between each pair among n nodes independently with probability p . We can think of $\mathcal{G}(n, p)$ equivalently as the model in which one starts with an n -node clique and, declares each edge to be live independently with probability p , and then considers the live-edge subgraph. The challenge in our case is that our graphs G are not necessarily cliques, or even close to being cliques, and relatively little is known about adapting results from $\mathcal{G}(n, p)$ to the case of arbitrary underlying base graphs [4]. Fortunately, however, we are able to prove a result that is strong enough for our purposes, adapting techniques for analyzing connected components in $\mathcal{G}(n, p)$ to the setting of live-edge subgraphs of arbitrary underlying graphs.

We begin with this part of the analysis, as follows.

(2.1) For all $\varepsilon > 0$ there exist constants $\alpha, \beta > 0$ such that the following holds. Let H be a graph in which each node has degree at least $r \geq \frac{1+\varepsilon}{p}$. Construct a random subgraph of H by declaring each edge to be “live” with probability p . Then for every node $i \in V$, the number of nodes reachable from i on live-edge paths is at least αr with probability at least β .

Proof. Let i be any node in H . We now describe a method for exploring the live edges outward from i , based on Karp’s analysis of random subgraphs of the bidirected complete graph [20] and Alon and Spencer’s analysis of infinite branching processes [5]. We first take all the nodes (if any) that i can reach via live edges and put them in a queue. We then repeatedly delete a node j from the queue and add to the queue all the nodes (if any) that j can reach via live edges, other than the ones already “discovered” (added to the queue) in previous iterations. Notice, crucially, that the outcome of the random live/blocked decision for each edge (j, j') is only examined once in this process, when one of nodes j or j' first comes to the front of the queue. Thus, we can assume that the live/blocked status of (j, j') is first determined at that moment.

For a small constant $\alpha > 0$, we say that this process *succeeds* if at least αr nodes are added to the queue before the queue ever becomes empty. If the process succeeds with probability at least β , for a constant $\beta > 0$, then our result follows.

Let Q_t be the number of nodes in the queue at the end of iteration t , where we define $Q_0 = 1$ to indicate that i starts in the queue. We have

$$Q_t = Q_{t-1} - 1 + X_t,$$

where the “-1” is because we delete a node j_t from the queue in iteration t (with $j_1 = i$), and X_t is a random variable equal to the number of not-yet-discovered nodes that j_t can reach via live edges. (This is where it is useful to assume that the live/blocked status of edges from j_t to not-yet-discovered nodes is only determined when j_t reaches the front of the queue.) Unrolling this recurrence, we have

$$Q_t = \left(\sum_{u=1}^t X_u \right) - t.$$

We are interested in showing that the probability of $Q_t > 0$ for all t from 1 until at least αr nodes have been discovered (added to the queue); in this case, the search for nodes using live-edge paths continues successfully for a sufficient number of steps, as required.

The expectation of X_t , prior to the point at which at least αr nodes have been discovered, can be determined as follows. The node j_t has degree at least r in H , and at most αr nodes have been discovered by the process thus far, so there are at least $(1 - \alpha)r$ edges emanating from j_t leading to not-yet-discovered nodes. We choose α small enough that $(1 - \alpha)r \geq \frac{1 + \varepsilon/2}{p}$; since each of these edges is live with probability p , we have $\mathbf{E}[X_t] \geq 1 + \varepsilon/2$. Thus, until αr nodes have been discovered, we can think of the queue length as a random walk on the integers with positive drift; as a result, there is a positive probability that the walk never returns to 0, which is the result we want.

We can briefly verify this in more detail for our particular case as follows. Let $S_t = \sum_{u=1}^t X_u$; by the Chernoff Bound, we have

$$\Pr[S_t \leq t] < \Pr[S_t \leq (1 - \varepsilon/4)\mathbf{E}[S_t]] < e^{-\frac{1}{2}\frac{\varepsilon^2}{16}t}.$$

Now, the sum $\sum_{t=1}^{\infty} e^{-\frac{1}{2}\frac{\varepsilon^2}{16}t}$ converges; we choose t_0 large enough that $\sum_{t=t_0}^{\infty} e^{-\frac{1}{2}\frac{\varepsilon^2}{16}t} < 1$. For p sufficiently small, there is a positive probability that X_1 , the number of nodes i can reach directly via live edges, is at least t_0 . It then follows that $S_t > t$ for all $t < t_0$. Finally, for all t we have $\Pr[S_t \leq t \mid X_1 \geq t_0] \leq \Pr[S_t \leq t] < e^{-\frac{1}{2}\frac{\varepsilon^2}{16}t}$; summing over t we obtain $\sum_{t=0}^{\infty} \Pr[S_t \leq t \mid X_1 \geq t_0] < 1$. ■

Next, we simply want to argue that if a node can reach many other nodes via live-edge paths with reasonably large probability, then it has a large probability of failing and hence a negative payoff. To do this, we first state a simple lemma about the union of many independent events (see the proof in the appendix), and then we use this to draw the resulting conclusion for a node's payoff.

(2.2) *Consider a collection of independent events $\mathcal{E}_1, \dots, \mathcal{E}_n$, each of probability $p > 0$. Then the probability of their union is at least $\min(\frac{1}{3}, \frac{2}{3}np)$.*

Now, for a node i , let the set of nodes it can reach on live-edge paths in G be called its *live component*, and let $r_i(G)$ be a random variable denoting the size of i 's live component.

(2.3) *For all $\gamma_0, \gamma_1 > 0$ there exist $\alpha, \delta > 0$ such that when $p, q \leq \alpha$ and Assumption $\mathcal{P}(\delta)$ holds, we have the following. If G is a graph with a node i for which $r_i(G) \geq \frac{\gamma_1}{p}$ with probability at least γ_0 , then the payoff of node i satisfies $\pi_i(G) < -bq$. (We note that the right-hand side is the payoff i would receive if it had no links).*

Proof. If i can reach at least $\gamma_1 p^{-1}$ nodes on live-edge paths, then by (2.2), the probability that it fails is at least $\min(\frac{1}{3}, \frac{2}{3}\gamma_1 p^{-1}q)$. Removing the conditioning on this event, the probability it fails is at least $\phi_i \geq \min(\frac{1}{3}\gamma_0, \frac{2}{3}\gamma_0\gamma_1 p^{-1}q)$. We also have $d_i \leq \Delta = c^* p^{-1}$.

If $\phi_i \geq \frac{1}{3}\gamma_0$, then

$$\begin{aligned}\pi_i &\leq ad_i - b\phi_i \leq ac^* p^{-1} - \frac{1}{3}b\gamma_0 \\ &\leq \delta bc^* - \frac{1}{3}b\gamma_0 = b(\delta c^* - \frac{1}{3}\gamma_0)\end{aligned}$$

where the last line is less than $-bq$ for δ sufficiently small and $q < \frac{1}{3}\gamma_0$.

If $\phi_i \geq \frac{2}{3}\gamma_0\gamma_1 p^{-1}q$, then defining $\gamma_2 = \frac{2}{3}\gamma_0\gamma_1$, we have

$$\begin{aligned}\pi_i &\leq ad_i - b\phi_i \\ &\leq ac^* p^{-1} - \gamma_2 bqp^{-1} \\ &= ac^* p^{-1} - (\gamma_2 p^{-1} - 1)bq - bq \\ &< \delta c^* b - (\gamma_2 p^{-1} - 1)bq - bq.\end{aligned}$$

This last line is less than $-bq$ provided that $1 + \delta c^* - \gamma_2 p^{-1} < 0$, which holds provided that p is sufficiently small relative to δ . ■

Finally, combining (2.1) with (2.3), we get an immediate consequence for the payoffs when all nodes have large degrees. The upper bound on min-welfare follows directly from this.

(2.4) *For all $\varepsilon > 0$, there exist $\alpha, \delta > 0$ such that when $p, q \leq \alpha$ and Assumption $\mathcal{P}(\delta)$ holds, we have the following. If each node has degree at least $\frac{(1+\varepsilon)}{p}$, then for each node i we have $\pi_i < -bq$.*

Proof. For each node $i \in V$, (2.1) implies that we have $r_i(G) \geq \frac{\gamma_1}{p}$ with probability at least γ_0 . It then follows from (2.3) that $\pi_i < -bq$. ■

(2.5) *For all $\varepsilon > 0$, there exist $\alpha, \delta > 0$ such that when $p, q \leq \alpha$ and Assumption $\mathcal{P}(\delta)$ holds, no graph can have min-welfare greater than $\frac{(1+\varepsilon)a}{p}$.*

Proof. Choose $\alpha, \delta > 0$ as in (2.4), and suppose by way of contradiction that there is a graph with min-welfare greater than $\frac{(1+\varepsilon)a}{p}$. It follows that every node i has degree greater than $\frac{1+\varepsilon}{p}$. But then by (2.4) we have $\pi_i < -bq$, contradicting the assumption that the min-welfare is greater than $\frac{(1+\varepsilon)a}{p}$. ■

3 Super-Critical Payoffs and Anonymous Markets

We now show that the upper bound in Section 2 can essentially be achieved, in an asymptotic sense, and also consider some of the structural implications of this fact.

To begin with, it is instructive to think about the analysis in Section 2 in terms of the random graph $\mathcal{G}(k, r)$.¹ One of the central facts about $\mathcal{G}(k, r)$ is that in a small window around probability $r = 1/k$, the expected size of the largest connected component jumps from a constant value to a constant fraction of k . This is the basic phase transition for $\mathcal{G}(k, r)$, and (2.1) in Section 2 is a reflection of this phase transition for an arbitrary underlying graph.

In order for a graph to achieve *super-critical payoffs* — those of the form $\frac{(1+\varepsilon)a}{p}$ for some $\varepsilon > 0$ — it must lie on the side of the phase transition where the live components are likely to be large, proportional to $1/p$. For this to be possible, it must cross the phase transition by little enough that these large components do not eliminate the payoff of the nodes. We now show how to do this, constructing a family of graphs built from disjoint cliques that achieve min-welfare of the form $\frac{(1+\varepsilon)a}{p}$.

Some Basic Facts about $\mathcal{G}(k, r)$. We begin by carefully stating some quantitative results about the phase transition in $\mathcal{G}(k, r)$ in a form that will be useful for the analysis. The proof of the first claim is given in the appendix.

(3.1) *Let C_i denote the component containing node i in $\mathcal{G}(k, r)$. If we fix some*

¹Since n and p are basic parameters in our model, we adopt the different variable names k and r in discussing $\mathcal{G}(k, r)$. Also, in keeping with standard terminology, we will often refer informally to $\mathcal{G}(k, r)$ as “a random graph,” as though it is a single graph rather than a distribution over graphs.

other node j and look at the event $j \in C_i$, then we have $\Pr [j \in C_i] = \frac{1}{k} \cdot \mathbf{E} [|C_i|] - \frac{1}{k}$.

Thus, looking at the probability a node belongs to i 's component is equivalent to looking at the expected size of i 's component.

The following pair of standard results describe the contrasting behavior of component sizes on opposite sides of $r = 1/k$.

(3.2) Fix $x < 1$, and consider the component of a given node i in $\mathcal{G}(k, r)$, where $kr = x$. Then for k sufficiently large, we have the following:

- (i) The probability that i 's component exceeds size c decreases exponentially in c .
- (ii) Consequently, the expected size of i 's component is bounded by a constant $c = c(x)$, independent of k , and the maximum size of any component in the graph is thus $O(\log k)$.

(3.3) There is an increasing function $\theta : [1, \infty] \rightarrow [0, 1]$ that is continuously differentiable on $(1, \infty)$ and continuously differentiable from the right at $x = 1$, with $\theta(1) = 0$ and $\theta'(1)$ a positive real number, such that the following holds. Fix $x > 1$ and $\varepsilon > 0$, and consider $\mathcal{G}(k, r)$, where $kr = x$. Then for $k = k(x, \varepsilon)$ sufficiently large, we have the following:

- (i) With probability $1 - \exp(-k)$, there is a component of size between $(1 - \varepsilon)\theta(x)k$ and $(1 + \varepsilon)\theta(x)k$.
- (ii) Conditioned on not belonging to the giant component in (i), the probability that a node i belongs to a component of size greater than c decreases exponentially in c .
- (iii) Consequently, the expected size of i 's component is between $(1 - \varepsilon)^2\theta(x)^2k$ and $(1 + \varepsilon)^2\theta(x)^2k + c$ for a constant $c = c(x)$.

Point (iii) follows from (i) and (ii) by considering that with probability $(1 \pm \varepsilon)\theta(x)$, node i belongs to a component of size $(1 \pm \varepsilon)\theta(x)k$, and with the remaining probability i belongs to a component of expected size at most c .

A Family of Graphs with Super-Critical Payoffs. For parameters k and s , let $F_s(k)$ denote the disjoint union of s cliques of size k . We will show that $F_s\left(\frac{1+\gamma}{p}\right)$, for arbitrary $s \geq 1$ and a small constant $\gamma > 0$, achieves super-critical node payoffs.

For our construction, we will focus on the special case $p = q$. A nice feature of this special case is that we can represent the spread of failures in $F_s(k)$ in the following equivalent way. We imagine a single “failure node” i^* associated with each clique, and attached to each real node in the clique, resulting in a clique on $k + 1$ nodes. There is a transmission probability p on the edges from i^* to each node in its clique, as there is on all other edges. In this view, a node i fails if it is in the same live-edge component as i^* ; in other words, the probability i fails is the probability it belongs to the same component as a given fixed node i^* in $\mathcal{G}(k + 1, r)$. By (3.1) we know this is

$$\frac{1}{k+1} \cdot \mathbf{E}[|C_i|] - \frac{1}{k+1},$$

where C_i denotes the live-edge component of i .

With $p = q$, we define σ to be the ratio $a/bp = a/bq$; by assumption $\mathcal{P}(\delta)$, we have $\sigma < \delta$, and we assume as usual that δ and p are sufficiently small. We let the number of nodes k in each clique be $(1 + \gamma)/p$ for a small value $\gamma > 0$ that we determine below.

First, (3.3)(iii) implies that the probability ϕ_i that i fails satisfies

$$(1 - \varepsilon_0)\theta(1 + \gamma)^2 \leq \phi_i \leq (1 + \varepsilon_0)\theta(1 + \gamma)^2$$

for a constant ε_0 that goes to 0 with p . Thus, the payoff to a node i is

$$\begin{aligned} \pi_i &\geq \frac{a(1 + \gamma)}{p} - \left(b + \frac{a(1 + \gamma)}{p}\right) (1 + \varepsilon_0)\theta(1 + \gamma)^2 \\ &= \frac{a(1 + \gamma)}{p} - b(1 + \sigma(1 + \gamma))(1 + \varepsilon_0)\theta(1 + \gamma)^2 \\ &\geq \frac{a(1 + \gamma)}{p} - 2b\theta(1 + \gamma)^2 \\ &= \frac{\sigma bp(1 + \gamma)}{p} - 2b\theta(1 + \gamma)^2 \\ &= b(\sigma(1 + \gamma) - 2\theta(1 + \gamma)^2). \end{aligned}$$

Now, let

$$h_0(x) = \sigma x - 2\theta(x)^2,$$

so that

$$\pi_i \geq bh_0(1 + \gamma).$$

We have

$$h'_0(x) = \sigma - 4\theta(x)\theta'(x).$$

Since $\theta(1) = 0$ and $\theta'(1)$ is a positive real number, we have $h'_0(1) = \sigma$, and hence the function $h_0(x)$ is strictly increasing over the interval $x \in [1, w_0]$ for a constant w_0 depending on σ . Since $h_0(1) = \sigma$, we have $h_0(w_0) = \sigma(1 + \sigma_0)$ for a constant $\sigma_0 > 0$ depending on σ .

Returning to the lower bound on π_i , we choose $\gamma = w_0 - 1$, and so

$$\pi_i \geq bh_0(w_0) = \frac{a}{\sigma p} \cdot \sigma(1 + \sigma_0) = \frac{a(1 + \sigma_0)}{p}.$$

Consequently, the payoff to each node exceeds $\frac{a}{p}$ by a multiplicative factor greater than 1 that depends on σ .

Comparison to an Anonymous Structure. The construction above achieves super-critical payoffs by allowing nodes to cluster into communities of an appropriate size, and thus to insulate themselves from failures originating in other communities. Drawing on a market motivation, it is interesting to ask whether super-critical payoffs can be achieved through structures that are based instead on *anonymous* interaction, where nodes can specify the number of partners they want to connect to, but have no control over who these partners are — the partners are chosen uniformly at random from the population. As we now show, in fact, anonymous interaction structures are not able to yield super-critical payoffs.

To define these anonymous structures precisely, we use the *configuration model* for random graphs [9, 10, 22]. Each of the n nodes is assigned k “half-edges”; these half-edges are then matched up uniformly at random into pairs, with each matching pair of half-edges forming an edge in the resulting random graph. Note that the pairing may cause two edges to go between the same pair of nodes, or for a node to form an edge that loops to itself; we remove these parallel edges and self-loops to obtain the final graph. Failures then propagate in this graph according to our model, spreading from root failures along live-edge paths.

With high probability, the local neighborhood of a node in this random graph will have a particularly simple structure, as follows. For node i , define $B(i, \ell)$ to

be the ball of radius ℓ centered at i , i.e. the induced subgraph of G on the set of all nodes reachable from i in ℓ or fewer hops. For fixed integers k, ℓ and any node i , the probability that $B(i, \ell)$ is a tree of depth ℓ and degree k (i.e. one whose internal nodes all have degree k and whose leaves are all at distance ℓ from the root) tends to 1 as $n \rightarrow \infty$.

For our analysis, we will therefore connect the propagation of failures in the configuration model to a related, simpler model based on an infinite k -regular tree. In particular, let $\mathcal{B}(k, r)$ denote the distribution over trees obtained by starting with an infinite k -regular tree and including each edge in the random tree with probability r . We now have a pair of results analogous to (3.2) and (3.3).

(3.4) *Let $x < 1$, and consider a tree generated from $\mathcal{B}(k, r)$ where $kr = x$.*

- (i) *The probability that the tree's size exceeds size c decreases exponentially in c .*
- (ii) *The expected size of the tree is bounded by a constant $c = c(x)$.*

(3.5) *There is an increasing function $\tau : [1, \infty] \rightarrow [0, 1]$ that is continuously differentiable on $(1, \infty)$ and continuously differentiable from the right at $x = 1$, with $\tau(1) = 0$ and $\tau'(1)$ a positive real number, such that the following holds. Consider a tree generated from $\mathcal{B}(k, r)$, and let $\psi_r(k)$ be the probability that it has an infinite node set.*

- (i) *If $kr > 1$, then $\psi_r(k) > \tau(kr)$.*
- (ii) *For all integers $c_0, c_1 > 1$ and $k \geq c_0 c_1 / r$, we have*

$$\psi_r(k) \geq 1 - (1 - \psi_r(c_0/r))^{c_1} > 1 - (1 - \tau(c_0))^{c_1}.$$

- (iii) *Conditioned on not having an infinite node set, the probability that the tree's size exceeds c decreases exponentially in c . Its expected size is thus bounded by a constant $c = c(x)$.*

(3.5)(iii) is a standard result; (i) and (ii) are formulated in ways that are adapted to our present purposes, and we give proofs of them in the appendix.

We now want to show that when each node forms k links in the anonymous structure, for any $k = \frac{1 + \beta}{p}$, the node payoffs can be at most a/p as $n \rightarrow \infty$. Clearly this is true for $\beta \leq 0$, so we consider the case of an arbitrary $\beta > 0$.

When the random graph G is sampled using the configuration model, for any node i the probability that the ball $B(i, \ell)$ is a tree of degree k and depth ℓ is $1 - o(1)$ as $n \rightarrow \infty$. Applying 3.5(i), the probability that i belongs to a live path of length ℓ is at least $\tau(1 + \beta) - o(1)$; for n sufficiently large, this probability is at least $\tau(1 + \beta/2)$. In the event that i belongs to a live path of length ℓ , it fails with probability at least $1 - (1 - q)^\ell$. By taking ℓ large enough, we may assume that $\tau(1 + \beta/2)(1 - (1 - q)^\ell) \geq \tau(1 + \beta/3)$ and thus node i fails with probability at least $\tau(1 + \beta/3)$.

Thus, if n is sufficiently large then we have

$$\begin{aligned} \pi_i &\leq \left(\frac{a(1 + \beta)}{p} \right) \left(1 - \tau \left(1 + \frac{\beta}{3} \right) \right) - b\tau \left(1 + \frac{\beta}{3} \right) \\ &= b\sigma(1 + \beta) \left(1 - \tau \left(1 + \frac{\beta}{3} \right) \right) - b\tau \left(1 + \frac{\beta}{3} \right). \end{aligned}$$

Let

$$h_1(x) = \sigma(1 + 3x) (1 - \tau(1 + x)) - \tau(1 + x),$$

so that $\pi_i \leq bh_1(\beta/3)$. By (3.5)(ii), we know that for $y \geq 4$, we have

$$\tau(y) \geq 1 - (1 - \tau(2))^{\lfloor y/2 \rfloor} \geq 1 - (1 - \tau(2))^{y/4}.$$

We can thus choose $w_1 \geq 4$ such that

$$\tau(y) \geq 1 - \frac{1}{1 + y}$$

for all $y \geq w_1$. If $1 + x \geq w_1$, we have

$$h_1(x) \leq \sigma(1 + 3x) \left(\frac{1}{2 + x} \right) - \frac{1 + x}{2 + x} < 0,$$

provided $\sigma < \frac{1}{3}$. Now, if $\sigma \leq \frac{1}{3} \sup_{y \in [1, w_1]} \tau'(y)$, then we have the following for all $x \in [0, w_1 - 1]$:

$$\begin{aligned} h_1'(x) &= 3\sigma(1 - \tau(1 + x)) - (\sigma + 3\sigma x + 1)\tau'(1 + x) \\ &\leq 3\sigma - \tau'(1 + x) \leq 0. \end{aligned}$$

Thus, for all $x \in [0, w_1 - 1]$, we have $h_1(x) \leq h_1(0) = \sigma$. Since we also have $h_1(x) < 0$ for $x \geq w_1 - 1$, it follows that $h_1(x) \leq \sigma$ for all $x \geq 0$.

Thus, for any $\beta > 0$, we have $\pi_i \leq b\sigma = \frac{a}{p}$ when each node forms $k = \frac{1 + \beta}{p}$ links. Since $\pi_i \leq \frac{a}{p}$ when nodes form at most $k \leq \frac{1}{p}$ links, it follows that for any constant c , if nodes form $\frac{c}{p}$ links then $\pi_i \leq \frac{a}{p}$ provided n is sufficiently large as a function of c .

Clustered vs. Anonymous Markets. It is instructive to consider why a union of disjoint cliques was able to achieve qualitatively higher payoffs than an anonymous interaction pattern. In particular, the nodes in the cliques we constructed are linking at a degree beyond the phase transition point, whereas attempting to do this in the anonymous structure has negative effects on the payoff.

A quantitative way to think about the contrast is to observe that in the union of cliques, the failure probability of a node i was approximately controlled by a conjunction of two events: i belonging to the giant component of the clique, and the “failure node” i^* also belonging to the giant component of the clique. As a result, the failure probability involves a term of the form $\theta(x)^2$, and this has a derivative of 0 at $x = 1$ — hence, it is safe to increase x a bit past 1 without blowing up the failure probability. On the other hand, in the anonymous structure, once i belongs to the giant component, it fails with overwhelming probability; thus, i ’s failure probability involves a term of the form $\tau(x)$, which has a strictly positive derivative at $x = 1$, and this makes it unprofitable to increase x even arbitrarily little past 1. This is the fundamental difference between the behavior of the two kinds of structures in the region just past the phase transition.

4 An Upper Bound on the Min-Welfare of Any Stable Network

We now show that any stable graph must have small min-welfare. (We defer the proof that stable graphs exist to the next section.) To upper-bound the min-welfare, we proceed as follows. First we show, in (4.1), that if two nodes i and j are not connected by an edge, and neither is at the maximum degree Δ , then at least one of them must have a large failure probability — this is what dissuades the other from forming the link.

It follows that in a stable network, all low-degree nodes of low failure probability must form a clique, since any unlinked pair of them would have an incentive

to connect. If the number of nodes n is sufficiently large, we can then find a node i that is far from this clique. Hence node i , and every node within a large number of steps of i , must have large degrees; we can thus apply an analogue of (2.1) to show i has a large failure probability, and this will conclude the proof.

(4.1) *Suppose Assumption $\mathcal{P}(\delta)$ holds. Let G be a stable graph, and let i and j be two nodes of G such that (i, j) is not an edge of G , and the degrees of i and j are each strictly less than Δ . Then we have $\max(\phi_i, \phi_j) \geq \frac{(1 - \delta)a}{(1 + \delta c^*)bp}$.*

Proof. Since the degrees of i and j are each strictly less than Δ , at least one of i or j does not have a strictly higher payoff if the edge (i, j) is included; let us assume it is node i . Thus, if G' denotes the graph G with the edge (i, j) included, then we have $\pi_i(G') \leq \pi_i(G)$.

We imagine evaluating failure in G' by first making all random root failure decisions and all random live/blocked decisions in G , then determining which additional nodes fail, and finally deciding whether the edge (i, j) is live or blocked and determining further failures. Let $\Phi_i(G)$ be the event that i fails in G before (i, j) is examined, and let $\mathcal{F}_{ij}(G)$ be the event that (i, j) is live and j fails in G . Then $\Phi_i(G') = \Phi_i(G) \cup \mathcal{F}_{ij}(G)$, so

$$\Pr[\Phi(G')] \leq \Pr[\Phi_i(G)] + \Pr[\mathcal{F}_{ij}(G)].$$

Since $\Pr[\mathcal{F}_{ij}(G)] = p\phi_j$, we have

$$\phi_i(G') - \phi_i(G) \leq p\phi_j(G).$$

Now,

$$\pi_i(G') = a(d_i + 1) - (ad_i + a + b)\phi_i(G'),$$

so the fact that $\pi_i(G') \leq \pi_i(G)$ implies that

$$a(d_i + 1) - (ad_i + a + b)\phi_i(G') \leq ad_i - (ad_i + b)\phi_i(G)$$

and hence

$$\begin{aligned} a &\leq (ad_i + a + b)(\phi_i(G') - \phi_i(G)) + a\phi_i(G) \\ &\leq (ad_i + a + b)p\phi_j(G) + a\phi_i(G) \\ &\leq (1 + \delta c^*)bp\phi_j(G) + \delta bp\phi_i(G), \end{aligned}$$

where the last line follows from the fact that $a < \delta bp$ and $d_i + 1 \leq c^*p^{-1}$.

Now, if $\phi_i(G) \geq \frac{a}{bp}$, we are done. Otherwise, we have

$$a \leq (1 + \delta c^*)bp\phi_j(G) + \delta a,$$

so

$$(1 - \delta)a \leq (1 + \delta c^*)bp\phi_j(G),$$

and hence

$$\phi_j(G) \geq \frac{(1 - \delta)a}{(1 + \delta c^*)bp}. \quad \blacksquare$$

Following our informal plan above, we note that a stable graph might have some low-degree nodes, so we require the following direct adaptation of (2.1), which applies to nodes that are far from all low-degree nodes.

(4.2) *For all $\varepsilon > 0$ there exist constants $\alpha, \beta > 0$ such that the following holds. Let H be a graph, and let A be the set of nodes of degree less than $\frac{1 + \varepsilon}{p}$. Let i be a node of distance greater than $\frac{1}{p}$ from A . Construct a random subgraph of H by declaring each edge to be “live” with probability p . Then the number of nodes reachable from i on live-edge paths is at least αp^{-1} with probability at least β .*

Proof. Consider the node-discovery process described in the proof of (2.1), starting from the node i , and recall that we declare it to *succeed* if it adds at least $\frac{\alpha}{p}$ nodes to the queue before it ever becomes empty, for the small constant $\alpha < 1$ used there. The event that the process succeeds depends only on the live/blocked decisions for nodes within distance $\frac{\alpha}{p}$ of i , and all such nodes have degrees at least $\frac{1 + \varepsilon}{p}$; hence, for this whole time we can apply the argument used in (2.1). \blacksquare

Finally, we conclude the proof strategy outlined at the beginning of the section, resulting in our upper bound.

(4.3) *Let $n > \Delta^{\Delta+2}$. For all $\varepsilon > 0$ there exist $\alpha, \delta > 0$ such that when $p, q \leq \alpha$ and Assumption $\mathcal{P}(\delta)$ holds, no stable graph can have min-welfare greater than $\frac{\varepsilon \alpha}{p}$.*

Proof. Suppose by way of contradiction that $G = (V, E)$ is a stable graph in which $\pi_i \geq \frac{\varepsilon a}{p}$ for all $i \in V$.

Let $A \subseteq V$ denote the set of all nodes i of G for which $d_i < \Delta$ and $\phi_i < \frac{(1 - \delta)a}{(1 + \delta c^*)bp}$. Since any node in A is able to form an additional edge, (4.1) implies that there must be an edge between each pair of nodes in A — in other words, A induces a clique in G .

Let $B \subseteq V$ denote the set of all nodes in G of degree equal to Δ . For any $i \in V - (A \cup B)$, we have $\phi_i \geq \frac{(1 - \delta)a}{(1 + \delta c^*)bp}$. Since $\pi_i \geq \frac{\varepsilon a}{p}$ by assumption, we have

$$\frac{\varepsilon a}{p} \leq ad_i - b\phi_i \leq ad_i - \frac{(1 - \delta)a}{(1 + \delta c^*)p}$$

and hence

$$d_i \geq \frac{\varepsilon}{p} + \frac{1 - \delta}{(1 + \delta c^*)p}.$$

For δ sufficiently small, the right-hand side of this inequality is at least $\frac{1 + \varepsilon_1}{p}$ for a constant $\varepsilon_1 > 0$. Choosing $\varepsilon_2 = \min(\varepsilon_1, c^* - 1)$, it follows that all nodes $i \in V - A$ have degree at least $\frac{1 + \varepsilon_2}{p}$.

Now, for any $j \in A$, there are at most $1 + \Delta + \Delta^2 + \dots + \Delta^{\Delta+1} < \Delta^{\Delta+2} < n$ nodes within distance $\Delta + 1$ of j , and hence within distance Δ of some node in A . Hence there is some node $i \in V$ at distance greater than $\Delta > p^{-1}$ from A . For this node i , (4.2) implies that $r_i \geq \frac{\gamma_1}{p}$ with probability at least γ_0 , for constants $\gamma_0, \gamma_1 > 0$. By (2.3), it follows that $\pi_i < -bq$, contradicting the assumption that the min-welfare of G is greater than $\frac{\varepsilon a}{p}$. ■

5 The Existence of Stable Networks

Finally, we show that there exist arbitrarily large stable networks. As with our constructions in Section 3, we will consider graphs that consist of disjoint cliques — graphs $F_s(k + 1)$ with $k = \frac{1 + \gamma}{p}$ for an appropriately chosen $\gamma > 0$.

The challenge is to find a k where the union of cliques is stable, and this requires some care for the following reason. Stability requires that no unlinked pair

of nodes wants to form an edge — this can be achieved by making k sufficiently large that creating a link between two cliques brings about too large an increase in failure probability to the nodes forming the link. Unfortunately, making k large also raises the failure probability of each node i based simply on its current set of edges — so we must not raise k so high that a node i wants to drop all its existing links. The crux of the problem is thus the following: is there a k that is large enough to discourage the formation of cross-clique links, but not so large that nodes will drop their current links? The main part of our analysis will be to show that such a k exists.

As in Section 3, we will consider the case in which $p = q$; defining σ to be the ratio $a/bp = a/bq$, we have $\sigma < \delta$, and we assume δ and p are sufficiently small. Due to space limitations, we prove the stability result in the appendix.

(5.1) *Given a, b, p, q as above, there exists $\gamma > 0$ such that with $k = \frac{1 + \gamma}{p}$, the union of cliques $F_s(k + 1)$ is stable.*

A Stable Graph with Unequal Clique Sizes. We observe that starting with a set of disjoint cliques $F_s(k + 1)$, we can create a different stable graph by adding one additional clique Γ of size $\ell < k + 1$ on a disjoint set of nodes. The size ℓ can be chosen in any way such that the payoffs of nodes in the clique Γ each exceed $-bq$. In this way, nodes in Γ will not want to drop their incident edges. Moreover, there is still no edge that can form so as to improve the payoffs of both its endpoints, since any edge involving a node i in Γ must have its other end at a node j in one of the cliques of size $k + 1$, in which case the argument for (5.1) shows that i would not want to form the link.

In particular, this means that we can take ℓ to be a clique yielding the maximum possible node payoff over all clique sizes, as in Section 3; this shows how certain nodes in a stable graph can have higher payoffs than others.

References

- [1] R. Albert, H. Jeong, A.-L. Barabási. Error and attack tolerance in complex networks. *Nature*, 406(2000).
- [2] F. Allen and A. Babus. Networks in finance. Tech. Report 08-07, Wharton Financial Institutions Ctr., 2008.

- [3] F. Allen and D. M. Gale. Financial contagion. *Journal of Political Economy*, 108(1):1–33, Feb. 2000.
- [4] N. Alon, I. Benjamini, and A. Stacey. Percolation on finite graphs and isoperimetric inequalities. *Annals of Probability*, 32:1727–1745, 2004.
- [5] N. Alon and J. Spencer. *The Probabilistic Method*. John Wiley & Sons, second edition, 2000.
- [6] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents.
- [7] V. Bala and S. Goyal. A non-cooperative model of network formation. *Econometrica*, 68:1181–1229, 2000.
- [8] M. A. Barnard. Needle sharing in context: Patterns of sharing among men and women injectors and HIV risks. *Addiction*, 88:805–812, 1993.
- [9] E. A. Bender and E. R. Canfield. The asymptotic number of labeled graphs with given degree sequences. *J. Combinatorial Theory (Series A)*, 24:296–307, 1978.
- [10] B. Bollobás. *Random Graphs*. Camb. U. Press, 2001.
- [11] R. J. Caballero and A. Simsek. Fire sales in a model of complexity. NBER Tech. Report 15479, 2009.
- [12] J. Corbo and D. C. Parkes. The price of selfish behavior in bilateral network formation. *Proc. 24th ACM Symp. on Principles of Distributed Computing*, 2005.
- [13] A. Fabrikant, A. Luthra, E. N. Maneva, C. H. Papadimitriou, S. Shenker. On a network creation game. *22nd ACM Symp. on Principles of Distributed Comp.*, 2003.
- [14] S. Goyal and A. Vigier. Robust networks, Apr. 2010. Working paper, Cambridge University.
- [15] A. Gutfraind. Mathematical terrorism, 2010. Ph.D. Thesis, Cornell University.

- [16] A. G. Haldane and R. M. May. Systemic risk in banking ecosystems. *Nature*, 469:351–355, 2011.
- [17] M. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.
- [18] M. Jackson, A. Wolinsky. A strategic model of social and economic networks. *J. Econ. Theory*, 71(1996).
- [19] J. Jacquez, C. P. Simon, J. Koopman, L. Sattenspiel, T. Perry. Modeling and analyzing HIV transmission. *Mathematical Biosciences*, 102(1988).
- [20] R. M. Karp. The transitive closure of a random digraph. *Random Structure and Algorithms*, 1(1):73–93, 1990.
- [21] É. Tardos, T. Wexler. Network formation games and the potential function method. *Algorithmic Game Theory* (Nisan, Roughgarden, Tardos, Vazirani, eds.), 2007.
- [22] N. Wormald. *Some problems in the enumeration of labelled graphs*. PhD thesis, Newcastle University, 1978.

6 Appendix

Proof of (2.2). If $p \geq \frac{1}{3}$ then the result follows immediately. Otherwise, if $np \leq \frac{2}{3}$, then we have

$$\begin{aligned}
 \Pr \left[\bigcup_{j=1}^n \mathcal{E}_j \right] &\geq \sum_{j=1}^n \Pr [\mathcal{E}_j] - \sum_{j,j'} \Pr [\mathcal{E}_j \cap \mathcal{E}_{j'}] \\
 &= np - \binom{n}{2} p^2 \\
 &\geq np - \frac{1}{2} (np)^2 \\
 &= np \left(1 - \frac{1}{2} np \right) \\
 &\geq \frac{2}{3} np.
 \end{aligned}$$

Otherwise, we can choose a subset S of $k \leq n$ of the events such that $\frac{2}{3} < kp \leq 1$. We have

$$\begin{aligned}
\Pr \left[\bigcup_{j \in S} \mathcal{E}_j \right] &\geq \sum_{j \in S} \Pr [\mathcal{E}_j] - \sum_{j, j' \in S} \Pr [\mathcal{E}_j \cap \mathcal{E}_{j'}] \\
&= kp - \binom{k}{2} p^2 \\
&\geq kp - \frac{1}{2} (kp)^2 \\
&= kp \left(1 - \frac{1}{2} kp \right) \\
&\geq \frac{2}{3} \cdot \frac{1}{2} = \frac{1}{3}.
\end{aligned}$$

■

Proof of (3.1).

$$\begin{aligned}
\Pr [j \in C_i] &= \sum_{s=1}^k \Pr [|C_i| = s] \cdot \Pr [j \in C_i \mid |C_i| = s] \\
&= \sum_{s=1}^k \frac{s-1}{k} \Pr [|C_i| = s] \\
&= \frac{1}{k} \sum_{s=1}^k (s-1) \Pr [|C_i| = s] \\
&= \frac{1}{k} \cdot \mathbf{E} [|C_i|] - \frac{1}{k}
\end{aligned}$$

■

Proof of (3.5). First we prove (i). The probability that the tree is infinite is the unique solution to $z = 1 - (1 - pz)^k$ in the interval $(0, 1)$. Define τ to be the unique solution to $\tau = 1 - e^{-x\tau}$ in the interval $(0, 1)$. Writing $f_0(v) = (1 - pv)^k$ and $f_1(v) = e^{-xv}$, we have

$$f_0(v) = (1 - pv)^k = (1 - pv)^{x/p} < e^{-xv} = f_1(v).$$

Thus, the curve $y = 1 - f_0(v)$ lies above the curve $y = 1 - f_1(v)$ on the interval $(0, 1)$, and so $y = 1 - f_0(v)$ intersects the line $y = v$ to the right of where

$y = 1 - f_1(v)$ intersects it. It follows that $z > \tau$, and hence we can take $\tau = \tau(x)$ as our function.

To prove (ii), consider $k' = \frac{c_0 c_1}{r}$ subtrees of the root in the complete k -ary tree (before edges are randomly included), and group them into c_1 blocks of $\frac{c_0}{r}$ subtrees each. For any block, if we consider just the root and the subtrees in a single block, the probability that the resulting random tree is infinite is at least $\psi_r(c_0/r)$ (since the root has this degree in the restricted tree, and the nodes in the subtrees have degree $k \geq c_0/r$). The tree is infinite if it is infinite in any of the blocks, and so the probability it is infinite is at least

$$1 - (1 - \psi_r(c_0/r))^{c_1} > 1 - (1 - \tau(c_0))^{c_1}$$

where the latter inequality follows directly from (i). ■

Proof of (5.1). For the analysis of the construction, we will work with the function $\theta(x)$ defined in (3.3), as well as the related function $\lambda(x) = x(1 - \theta(x)^2)$. Observe that $\lambda(1) = 1$, since $\theta(1) = 0$. Taking derivatives, we have

$$\lambda'(x) = (1 - \theta(x)^2) - 2x\theta(x)\theta'(x),$$

and hence $\lambda'(1) = 1$. Thus we have

(6.1) *For some constant $w > 1$, the function $\lambda(x)$ is strictly increasing on the closed interval $[1, w]$.*

As in Section 3, we analyze the failure process by attaching a single “failure node” i^* to each clique. The probability ϕ_i that node i fails is the probability that i belongs to the same live-edge component as i^* in the $(k + 2)$ -node clique where i^* is added to i ’s clique. The payoff to node i is

$$\pi_i = ak - (ak + b)\phi_i.$$

If i drops all its edges, it receives a payoff of $-bq < 0$. If i forms an edge to a node j in another clique, it receives an added benefit of a , and incurs an increased expected loss of at least

$$(ak + b)p\phi_i(1 - \phi_i).$$

There are four terms here; the second and third represent the chance that j ’s failure (which is $\phi_j = \phi_i$ by symmetry) spreads to i , and the fourth term represents the

fact that this only matters if i had not already failed in its own clique. In more detail: the payoff to node i before the addition of this edge is $ak - (ak + b)\phi_i$, and afterward it is

$$a(k + 1) - (ak + a + b)(\phi_i + p\phi_i - p\phi_i^2),$$

so the change in payoff is less than $a - (ak + b)p\phi_i(1 - \phi_i)$.

Now, what is ϕ_i ? By (3.1) and (3.3), we have

$$(1 - \varepsilon_1)^2\theta(p(k + 2))^2 \leq \phi_i \leq (1 + \varepsilon_1)^2\theta(p(k + 2))^2 + c_1p$$

for a constant ε_1 that goes to zero as p does. By choosing a slightly larger ε_2 , and using the fact that $\theta(\cdot)$ has a bounded first derivative, we have

$$(1 - \varepsilon_2)\theta(1 + \gamma)^2 \leq \phi_i \leq (1 + \varepsilon_2)\theta(1 + \gamma)^2,$$

with ε_2 going to zero as p does.

In the expression $\phi_i(1 - \phi_i)$, provided the upper bound $(1 + \varepsilon_2)\theta(1 + \gamma)^2 \leq \frac{1}{2}$, we have

$$\begin{aligned} \phi_i(1 - \phi_i) &\geq (1 - \varepsilon_2)\theta(1 + \gamma)^2(1 - (1 - \varepsilon_2)\theta(1 + \gamma)^2) \\ &\geq (1 - \varepsilon_2)\theta(1 + \gamma)^2(1 - \theta(1 + \gamma)^2). \end{aligned}$$

Since $a = \sigma bp$ and $k = (1 + \gamma)/p$, if we write $\sigma_1 = \sigma(1 + \gamma)$, then we have $ak = \sigma(1 + \gamma)b = \sigma_1 b$. Now we have

$$\begin{aligned} (ak + b)p\phi_i(1 - \phi_i) &= b(1 + \sigma_1)(1 + \gamma)\phi_i(1 - \phi_i) \\ &\geq b(1 + \sigma_1)(1 + \gamma)(1 - \varepsilon_2)\theta(1 + \gamma)^2(1 - \theta(1 + \gamma)^2) \\ &= b(1 + \sigma_1)(1 - \varepsilon_2)\lambda(1 + \gamma)\theta(1 + \gamma)^2 \\ &\triangleq f_1(\gamma), \end{aligned}$$

where the last line is taken as the definition of $f_1(\gamma)$. Observe that $f_1(0) = 0$, and by (6.1), there is an $x_1 < 1$ such that the function $f_1(x)$ is strictly increasing for x in a closed interval $[0, x_1]$.

We also have

$$\begin{aligned} (ak + b)\phi_i &\leq b(1 + \sigma_1)(1 + \varepsilon_2)\theta(1 + \gamma)^2 \\ &\triangleq f_0(\gamma), \end{aligned}$$

where once again the last line is taken as the definition of $f_0(\gamma)$. We see that $f_0(x)$ is also strictly increasing in $[0, x_1]$ (and beyond this interval as well).

Now, since $\lambda(1) = 1$ and $\lambda(\cdot)$ is monotone increasing on $[0, x_1]$, for any small enough $\varepsilon_2 > 0$, there is a unique $x_0 < x_1$ such that

$$\lambda(1 + x_0) = \frac{1 + \varepsilon_2}{1 - \varepsilon_2}.$$

Moreover, $f_1(x) > f_0(x)$ for all $x \in (x_0, x_1]$, and the value of x_0 goes to 0 as ε_2 goes to 0. Also, we observe that for $\gamma \in (x_0, x_1]$, we have

$$f_1(\gamma) > f_0(\gamma) \geq b\theta(1 + \gamma)^2.$$

Now, we choose σ small enough that $2\sigma < \theta(1 + x_1)^2$. We then choose ε_2 small enough (by choosing p small enough) so that $f_1(x_0) = f_0(x_0) < b\sigma$. Finally, let $g(\gamma) = ak = (1 + \gamma)b\sigma$. Since $b\sigma < g(\gamma) < 2b\sigma$ for all $\gamma \in (x_0, x_1]$, it follows that $f_0(x_0) < g(x_0)$ but $g(x_1) < f_1(x_1)$. Therefore, since $f_0(\cdot)$ and $f_1(\cdot)$ are continuous functions, there exist $\gamma^*, \gamma^{**} \in (x_0, x_1]$ for which $g(\gamma^*) = f_1(\gamma^*)$ and $g(\gamma^{**}) \geq f_0(\gamma^{**})$, with $\gamma^* < \gamma^{**}$.

We choose any $\gamma \in [\gamma^*, \gamma^{**}]$ as the value of γ we use to define k . With this value of k , the payoff i receives from keeping all its edges is

$$\pi_i = ak - (ak + b)\phi_i \geq g(\gamma) - f_0(\gamma) \geq 0,$$

and hence i prefers to keep its edges rather than deleting all of them. The change in payoff i receives from linking to a node j in a different clique is less than

$$\begin{aligned} & k^{-1} (ak - (ak + b)pk\phi_i(1 - \phi_i)) \\ & \leq k^{-1} (ak - b(1 + \sigma_1)(1 - \varepsilon_2)\lambda(1 + \gamma)\theta(1 + \gamma)^2) \\ & = k^{-1} (g(\gamma) - f_1(\gamma)) \\ & \leq 0, \end{aligned}$$

and hence i will not form this link. Thus, the graph $F_s(k + 1)$ is stable. ■