Cooperation in Anonymous Dynamic Social Networks*

Preliminary Working Paper

Brendan Lucier† Brian W. Rogers‡ Nicole Immorlica§

April 15, 2013

Abstract

We study the extent to which cooperative behavior can be sustained in large, anonymous, evolving social networks. Individuals strategically form relationships under a social matching protocol and engage in prisoner's dilemma interactions with their partners. We characterize a class of equilibria that support cooperation as a stationary outcome. When cooperation is possible, its level is uniquely determined. While neither community enforcement nor contagion mechanisms have force in our setting, the endogenous dynamics of the social network imply that cooperation allows an individual to gradually accumulate a large network of profitable interactions, while defection results in social marginalization. Even as players become perfectly patient, equilibrium allows for full cooperation, only autarky, or the coexistence of cooperation and defection, depending on payoffs. Smaller levels of cooperation can be sustained by a form of exclusivity among cooperators.

*We thank Simon Board, Wiola Dziuda, Matthew Jackson, Christoph Kuzmics, David Miller, Paolo Pin, Marzena Rostek, Bill Sandholm, Joel Sobel, Joel Watson and Simon Weidenholzer for helpful comments and conversations.

†University of Toronto, Toronto, ON Canada, blucier@cs.toronto.edu

‡MEDS, Kellogg and NICO, Northwestern University, Evanston, IL USA, b-rogers@kellogg.northwestern.edu

§Northwestern University, Evanston, IL USA. nicimm@gmail.com

1 Introduction

In the modern world, many social interactions and economic transactions are mediated through large evolving networks of agents. The size of these systems along with the transient nature of both individual membership and the relationships amongst them often confer a measure of anonymity to participants. In online contexts, there are technological aspects that reinforce anonymity, in that it is typically possible to create a new identity at any point in time and thereby mimic a new entrant. In these settings, the scope for punishing uncooperative behavior and, hence, for sustaining efficient outcomes, is potentially limited. Nevertheless, many systems that are largely anonymous and in which agents change partners over time are characterized by a high, though often less than universal, level of cooperation. We find that such an outcome can be explained as equilibrium behavior under a simple network formation model.

One important aspect of social and economic networks is that an agent's strategic behavior within the context of a particular relationship is influenced by the behavior from other relationships. This could in fact be taken as a defining characteristic of any interesting network, as in the absence of some form of strategic interdependence, one is essentially engaged in a set of separable relationships and the network, per se, has no influence on behavior.

We consider a strong form of this dependence in which at any given time each agent takes the same action with each of his neighbors. There are two reasons for adopting this approach. First, there are applications in which technological constraints limit an agent's ability to behave differently with different partners, or in which the chosen action represents a general characteristic of an agent that is not relationship-specific. This is the case if, for example, an agent must invest in a quality for each period and use that quality for each transaction. Second, if there is even a small amount of local information flow in the network, such that a defection with one partner would be observed by the agent's other partners, then there will exist equilibria with the behavior that we describe here even if there are no restrictions at all on the profile of actions an individual chooses at a given moment in time.[1]

We now describe the main elements of our framework. Agents enter the system over

[1] One could view this as an implementation of "local community enforcement" along network connections, in the spirit of [12, 6], which we discuss in more detail below.

time and have finite lives. All strategic interactions are bilateral and described by a prisoners' dilemma. A random matching process presents agents with opportunities to form new relationships. In every period, each agent chooses a behavior, cooperation or defection, and receives the sum of payoffs from the corresponding stage games with each of its current partners. After every period, each agent has the opportunity to sever any of its relationships.

Consider, for example, an online community in which agents seek partners with whom to profitably interact, such as for trading goods or engaging in joint production. At any point in time agents can choose to conduct honest business (cooperate) or to cheat their partners for a gain (defect). The discretion to sever a relationship has an important impact on behavior. In the model it plays a key role by providing a mechanism with which to threaten punishment for uncooperative behavior. In fact, because of anonymity, this is the only effective mechanism for punishment since there is no scope for future partners to punish a defecting agent.[2] In addition to severing a link to a defecting agent, one might wish to broadcast the agent's defection so as to enable further punishment. But the agent who defected does not have an identity that can be tracked by future partners, and so bears no negative consequence of his defection beyond the potential loss of his current relationships.[3]

Interestingly, [18] provide empirical evidence based on laboratory experiments that the ability to endogenously determine one's partners increases observed cooperation levels in repeated prisoner's dilmmea interactions, while [1] show that cheaters in an online game are punished by an increased rate of losing partners. These effects constitue crucial elements of the model we study.

There are two essential properties of the matching process that drive our results. The first is that it takes (valuable) time to search for partners with whom to form *profitable* relationships. In an environment characterized by anonymity and the ability to sever relationships, this is generally a necessary feature to provide any possibility of cooperative behavior. The intuition is that, because defection can be punished only by the severance of relationships,

[2] [11] define and study social games that changes equilibrium outcomes by permitting players to choose with whom they interact. This produces very different insights than our work, since here players are matched through a random process. [4] identify collaborative equilibria as a function of the social network that describes interactions.

[3] Because the defecting agent also has the option of severing the relationship, one can not threaten contigent play that reduces the defecting agent's continuation payoff from the relationship below its outside option.

in order for this punishment to bite, it must be that future relationships are sacrificed as a result of defecting.

In light of this observation, our mechanism for sustaining cooperation can be interpreted as a particular implementation of the notion of social capital, here taken to mean an agent's accumulated network of cooperative partners with whom he is connected.[4] Indeed, the notion that cooperative behavior is determined in large part by social pressures has been studied at length in the sociology literature, including, for example, [13]. In our setting, the reason to cooperate comes from the fact that, through cooperation, one can gradually build up a large social network consisting of other cooperating agents.

The second property is that there is a limit to the number of relationships that an agent can maintain. In combination with the first property, this dictates that the maginal returns to cooperating to be decreasing in the aggregate level of cooperation in the system. This concavity is important because it allows for the possibility that cooperating and defecting agents will co-exist in equilibrium. The coexistence result obtains for a wide range of parameters and we view it as having descriptive value.

Moreover, the constraint on the number of relationships has strategic importance in that it forces an agent to trade off the continuation value of each relationship against the outside option of initiating a new relationship. The value of the outside option is dictated by the aggregate system dynamics, and so provides a link between overall behavior in society and the incentives that govern behavior in a particular relationship.

Because of this strategic effect, our work is related to a branch of literature that studies repeated *bilateral* interactions when relationships can be endogenously terminated. The central idea in this literature is that the threat of severance disciplines behavior because being re-matched entails a cost. The cost can come in many forms, such as being cast into a matching market with frictions, as in the pioneering work of [19], having to start a new relationship that requires a specific investment [17], or having to start with small stakes in a new relationship [23, 24].[5] New relationships may also entail a phase of gradual trust

[4]See [21] for an overview of the large literature regarding the concept of social capital. With a related motivation, but very different analysis, Vega-Redondo [22] studies social capital in a stochastically evolving network.

[5][16] and [3] contribute to the search literature in the context of relational contracts.

building, as in [8] and [5], or the payment of a bond, as in [14].[6]

A crucial element common to all of this research is that an agent is involved in at most one relationship at any given time. All of the equilibria that are studied by this literature use strategies with the property that actions evolve in a non-trivial way throughout the course of a given relationship. In a networked society in which an agent continues to form new relationships while maintaining older relationships, these constructions have no natural analogue.[7] We focus instead on a particularly simple and intuitive kind of behavior that we refer to as *consistent*. A consistent agent chooses to either perpetually defect or to perpetually cooperate.[8] We exhibit equilibria in which all agents take consistent actions. In light of this behavior, optimal decisions regarding the formation and severance of relationships become easy to describe, which allows us to precisely pin down the nature of the co-evolution of the network and the behavior in repeated interactions occurring on the network. In particular, a relationship is severed when, and only when, a defection is observed. Such a social norm is very natural: defection is not tolerated, and cooperation is met with the opportunity for future interactions.

Assuming first that agents use consistent strategies, we provide a characterization of stable stationary equilibria (SSE). The first message is that the payoffs of the prisoner's dilemma have an important impact on the sustainable level of cooperations. Under adverse conditions, no cooperation can be sustained, even as players become perfectly patient. Otherwise, when agents are patient enough, there is a unique SSE that supports cooperation, and it is such that either there is universal cooperation, or cooperators and defectors co-exist at a specific ratio.[9] In this sense our model provides an explanation of which societies permit universal cooperation, and which societies will necessarily be subject to a fringe of cheating behavior. The model allows for simple comparative statics on the equilibrium level of defection. This provides a first step towards assessing which kinds of policies can most effectively improve

[6] [20] provides an earlier analysis of a related game.

[7] In a very different framework, [9] also leverage the idea that cooperation can be enforced through the use of multiple relationships, and so make a related connection to social capital. [15] and [2] investigate relational contracts between a principle and many agents engaged in a repeated game.

[8] In a model where agents have only one partner at a time [7] examines strategy profiles that involve agents taking different strategies, similar to our roles of cooperation and defection, and partially characterize equilibria with trust building via an evolutionary approach.

[9] The autarky equilibrium always exists and is stable. When there is a positive SSE, there also exist either one or three unstable stationary equilibria.

the welfare of the system.

The way in which cooperators manage their relationships takes one of two forms. First, it may be that they always accept new relationships when the opportunity arises. This case obtains when the pool of agents searching for a match is sufficiently cooperative. Otherwise, they accept new relationships only with some probability $p < 1$, thereby imposing a barrier on society to obtaining connections with cooperators, resulting in a form of *exclusivity*. This exclusivity implies that some agents will fail to find partners when given the opportunity to match. It can thus be interpreted in terms of an endogenous friction operating on the matching process. This friction is necessary to achieve cooperation under some parameters. While exclusivity is costly to all agents, it is effective because, under the right circumstanced, it decreases the expected returns to defection by more than it decreases the returns to cooperation.

Having characterized SSE under consistent behavior, our seecond main result shows that under an appropriate condition on parameters, the outcomes we have identified describe play along an equilibrium path without imposing consistency on action choices. The condition, which requires that the temptation payoff to defect be small enough relative to the loss from being defected on, ensures that perpetual cooperation is sequentially rational at any possible history. This demonstrates that consistent behavior can be thought of as self-enforcing when paired with the social norm of severing a link when, and only when, a defection is observed.

Finally, we examine behavior when this condition is not met. The only kind of profitable deviation from the consistent strategy profile involves a cooperator defecting under a particular circumstance: when he has very little social capital, and therefore little to sacrifice in terms of future links with cooperators. We show that as the network becomes dense, this circumstance becomes increasingly rare due to a law of large numbers argument, and as a result these strategies form an epsilon equilibrium that generate essentially the same level of cooperation that we characterize under consistent strategies. In this sense, the description of consistent equilibria remains a valid description of outcomes for the model.

The remainder of the paper is organized as follows. The model is described in Section 2. Section 3 characterizes stationary equilibrium outcomes under consistent strategies. Section 4 shows that consistency is a self-enforcing norm under the appropriate condition, while Section 5 describes outcomes that involve a minimal failure of consistency. We conclude

and provide comments for further research in Section 6. An Appendix contains a formal development of the model and proofs.

2 A model of strategic interactions in a social network

For ease of exposition, we describe the main elements of the model in this section, relegating the full development of some technical aspects to the Appendix.

All strategic interactions are governed by a prisoner's dilemma with the following payoff matrix.

	C	D
C	1,1	-b,1+a
D	1+a,-b	0,0

We take $a, b > 0$ and $a - b < 1$ so that, while mutual cooperation is the uniquely efficient outcome, defection is strictly dominant.

There is a continuum of agents, which we associate with points from the unit interval $N = [0, 1]$.[10] Agents interact repeatedly on an evolving directed network. Time is discrete. At each date each agent independently dies with a given probability $1 - \delta$, in which case it is replaced by a new agent.[11] We speak of the age of an agent i, $t(i)$, as the number of dates since its birth, so that $t(i) = 0$ in the periods when i is born. Each agent i chooses an action $\alpha_i^{t(i)} \in \{C, D\}$ at each date $t(i)$ of its life. Agents commonly observe the aggregate proportion, q, of C behavior in the population after each date, and the aggregate proportion p of proposed inlinks that were accepted by cooperators[12].

Every agent is able to sponsor a number $K \geq 1$ of connections to other agents. Thus an agent is generally involved both in relationships that it sponsors (outlinks) and also in relationships sponsored by others (inlinks), resulting in a directed graph of interactions. When a connection is proposed, the partner is chosen uniformly at random from the population,

[10] We work with a continuum of agents so that players are "atomless", i.e., no individual can unilaterally affect aggregate behavior. This is a good approximation for large societies, as one then expects a player to ignore the marginal effect of her behavior aggregate play.

[11] The conclusion that at each date, a proportion δ of the population survives with probability one relies on an exact law of large numbers for a continuum of random variables. See, e.g., Judd (1985).

[12] Common knowledge of these summary statistics is useful for our notion of equilibrium stability, in which agents respond to an exogenous shock to the state, as discussed in Section 3.3. In particular, this information is not necessary to support equilibria.

and the connection is then accepted or rejected by the chosen partner. Once accepted, each connection persists to the subsequent date unless one of the partners dies or chooses to sever the connection. When a connection is broken, the agent who sponsored it, provided he survives, is able to re-match with another agent, chosen uniformly at random, at the next date.

Notice that there is an implicit bound on the expected number of inlinks an agent will ever receive over the course of his life, due to the fact that every inlink corresponds to the outlink of some other agent. The fact that we explicitly bound the number of outlinks, while leaving the bound on inlinks implicit does not drive the results. Similar results would obtain if one instead bounds the total number of (in- and out-) links an agent is able to maintin.

At each date an agent receives a payoff equal to the sum of the outcomes of the stage game played with each of his (in and out) partners, according to the chosen actions of the two agents and the payoff matrix given above. Agents seek to maximize the present value of expected lifetime payoffs.

To summarize, each time period proceeds according to the following order of events:

1. New agents are born.
2. (p, q) from the previous data is publicly observed
3. Actions are chosen.
4. Outlinks are proposed to other agents.
5. Potential inlinks are accepted or rejected.
6. The stage game is played and payoffs are realized.
7. Agents sever any links that they choose to.
8. Death occurs.

3 Consistent Behavior and stationary outcomes

An agent's strategy specifies at each date, as a function of everything the agent has observed, whether to cooperate or defect, how many links to propose, which proposed inlinks to accept, and which existing links to sever. A formal development is contained in the Appendix. Throughout, we maintain two assumptions on strategies. First, at the individual level, we assume a weak form of stationarity, in that agents do not condition their plan of action on

a common labeling of time. In other words, while each agent is aware of his age and the history he has observed, he does not use any universal description of time. Second, at the collective level, we assume that strategy profiles are symmetric.[13] However, as we will see, agents will indeed take heterogeneous roles along a given path of play.

Even with these two assumptions in place, the analysis involves many agents, birth and death, an endogenously evolving network of relationships, and histories that are to large extent privately observed. In order to gain traction on studying equiibrium outcomes, we begin the analysis by considering a setting in which agents adopt a particularly simple behavioral rule. Namely, agents are assumed to be consistent in their choice of action, cooperation or defection, over the course of their lives.

Property **1** *Consistency: An action from $\{C, D\}$ is chosen at birth (possibly mixing). At all future dates, the agent takes same the action it played at the previous date. That is, a strategy for i is consistent if for every $t(i) \geq 0$, $a_i^{t(i)+1} = a_i^{t(i)}$.*

Consistency allows us to speak of society as consisting of "cooperators" and "defectors". This behavior is plausible as well as simple. Moreover, cooperation will not require the use of elaborate punishment phases in equilibrium construction.

Certainly, though, consistency limits the complexity of strategic interactions in an important way. In particular, it prohibits strategies that allow an agent to cooperate until a history with a certain property is reached, and then defect. Even though consistency rules out many of the standard constructions that permit cooperation, it turns out to be a significantly milder condition than it may first appear. We formalize this assertion in Section 4, but the intuition is as follows. If cooperation at a given round is part of an optimal strategy, then it is because the increased access to it provides to relationships tomorrow is enough to forgo the temptation payoff. But if that is true today, then tomorrow the same comparison is likely to hold true again. To the extent that connection is valid, optimality of consistent cooperation is not a much stronger requirement than the optimality of the initial act of cooperation.

[13] Notice that we thus are not concerned with measurability assumptions on strategy profiles with a continuum of agents.

We will first demonstrate that interesting aggregate outcomes obtain under consistent strategies. Second, and more importantly, the analysis of consistent strategies is a preliminary step to studying equilibrium outcomes more generally. We will argue below that consistent strategies are, to large extent, self-enforcing, and describe much about equilibrium behavior in our framework even when inconsistent behavior is permitted.

Under consistent strategies, it is straightforward to describe the optimal management of links. First, if a defection is ever observed, the best response of the defector's partner is to sever the link. We refer to this as *unforgiving* behavior. Second, when cooperation is observed, the best response is to maintain the link (provided both agents survive), which we refer to as *trusting* behavior. Third, it remains to be shown under which conditions a cooperator should accept a proposed inlink.[14] This will be determined by the probability that a proposed link comes from a cooperator. Notice that this probability is less than q in equilibrium since, as just argued, defectors lose their links at every period, and have a dominant strategy to propose all possible links, while cooperators generally maintain some links from previous periods, and so search less. Because of anonymity, all proposed inlinks are ex ante equivalent from the perspective of the agent receiving the proposal. It is therefore essentially without loss to assume that a cooperator accepts each proposed inlink independently with a certain probability p. Finally, for any choice of p and q, we must confirm that it is indeed rational for cooperators to send outlinks. When $p = 1$, the chance of finding a cooperator is q, and so this condition is implied by the willingness to accept inlinks, but otherwise the condition is independent.

In light of these observations, behavior can be completely described by a function $\phi : [0,1]^2 \to [0,1]^2$ with the interpretation that $\phi(p,q)$ specifies the probability that an agent chooses C at its birth and the probability with which each inlink is accepted in the event the agent becomes a cooperator, upon observing the state (p,q).

3.1 Simple stationary equilibria

We are interested in determining when a particular level of cooperation q can be sustained as a stationary outcome of the system under consistent strategies. Note, however, that a given pair (p,q) does not, by itself, capture all of the payoff-relevant aspects of the system, even

[14] Defectors, of course, have a dominant strategy to accept every inlink

in expectation. Other factors, such as the amount of search by cooperators and defectors, which depends on the age distribution of those behaviors, impact expected payoffs. This motivates us to describe the state of the system under stationary behavior.

Definition **1** *The* steady-state at (p, q), $L_{(p,q)}$, *is the limiting distribution over graphs that obtains when all agents have been applying the consistent, unforgiving, and trusting strategy described by* (p, q) *for* t *periods, as* $t \to \infty$.

From the steady-state we can extract all payoff-relevant quantities for expected utilities, such as the age distribution of cooperators and defectors, along with the expected amount of search agents of either type are conducting at each age.

For a steady-state to be supported as an equilibrium outcome, it is necessary and sufficient that the strategy $\phi(p, q) = (p, q)$ be optimal when the system is in state $L_{(p,q)}$. In this case, the fact that all agents apply ϕ implies that the system remains in steady-state $L_{(p,q)}$. This is extremely useful for the analysis, since it is only under the steady-state assumption that we are able to derive expressions for expected utilities. We can now define equilibrium under consistent strategies.

Definition **2** *A pair* (p, q) *is a* simple stationary equilibrium (SSE) *if, given that the system is in* $L_{(p,q)}$ *at all times, the application of a strategy that chooses cooperation with probability* q *at birth, and where cooperators accept each inlink independently with probability* p*, is optimal in the space of consistent strategies.*

3.2 Expected Utilities

We now derive the expected utilities associated with the (consistent) choices of cooperation and defection at an agent's birth. These utilities depend on the model's parameters, (a, b, δ). They depend as well on the proportion of cooperative agents in society, q, and the rate of inlink acceptance, p. Since we are interested in simple stationary equilibria, we work under the assumption that the system is in state $L_{(p,q)}$ and remains so over the agent's lifetime. If (p, q) is to be a simple stationary equilibrium it is rational for agents to compute their utilities under the expectation that the system remains in $L_{(p,q)}$.

The main task in computing expected utilities is to keep track of the expected number of inlinks and outlinks between agents of different behaviors, C and D, as a function of age. Define $n_{XY}^{Out}(s)$ as the expected number of outlinks from an agent of type X at age s to agents of type Y, $X, Y \in \{C, D\}$. The expected number of links from a cooperator of age s to other cooperators can be computed recursively according to

$$n_{CC}^{Out}(s) = \delta n_{CC}^{Out}(s-1) + pq(K - \delta n_{CC}^{Out}(s-1)).$$

The first term retains the existing links with cooperators who remain alive, while the second term takes all links from the previous period that were broken (due to death or defection) and re-matches them, obtaining a fraction q of new cooperators, p of whom accept the link. Setting $n_{CC}^{Out}(-1) = 0$ and solving produces

$$n_{CC}^{Out}(s) = pqK \left(\frac{1 - (\delta(1-pq))^{s+1}}{1 - \delta(1-pq)} \right).$$

The number of links from a cooperator of age s to defectors can then be computed according to $n_{CD}^{Out}(s) = (1-q)(K - \delta n_{CC}^{Out}(s-1))$, since each link that is proposed at age s matches with a defector with probability $1 - q$s. For defectors the case is much simpler, as the property of unforgiving behavior implies that age dependency is trivial. We have $n_{DC}^{Out}(s) = pqK$ and $n_{DD}^{Out}(s) = (1-q)K$.

We turn now to the expected number of inlinks from both types of nodes as a function of age. To do so, we first compute the number of inlinks an agent expects to receive from agents of either behavior at each date. These are time-independent rates in steady-state, and from them the evolution of inlinks is easy to derive. The probability that a randomly selected node is age s is $f(s) = (1-\delta)\delta^s$. In steady-state, $f(s)$ also defines the age distribution of cooperators and the age distribution of defectors. Then, the expected number of inlinks an agent will receive from cooperators and defectors at each date are, respectively,

$$\begin{aligned} r_C &= q \sum_{s=0}^{\infty} f(s) \left(K - \delta n_{CC}^{Out}(s-1) \right) = qK \frac{(1-\delta^2)}{1 - \delta^2(1-pq)}, \\ r_D &= (1-q)K. \end{aligned}$$

Notice that the calculation of r_C requires the assumption that the system is in a steady-state, since it presumes that for every age s, the proportion of age-s agents that cooperate is q.

(The calculation for r_D, on the other hand, is valid for any state consistent with q since the number of outlinks sent by a defector is independent of age.)

Define $n^{In}_{XY}(s)$ as the expected number of inlinks an agent of type X at age s has from agents of type Y, $X, Y \in \{C, D\}$. For CC links, we have the recursive relationship

$$n^{In}_{CC}(s) = \delta n^{In}_{CC}(s-1) + r_C.$$

Setting $n^{In}_{CC}(-1) = 0$ and solving produces

$$n^{In}_{CC}(s) = r_C \frac{1-\delta^{s+1}}{1-\delta}.$$

The remaining calculations are straightforward since they all involve defectors whose links are re-set every period. We have $n^{In}_{CD}(s) = n^{In}_{DD}(s) = r_D$ and $n^{In}_{DC}(s) = r_C$.

Finally, we can now define the expected lifetime utility of consistently cooperating and consistently defecting. To that end we compute the expected payoff at a particular age s by summing the payoffs over the expected set of connections. We have

$$\begin{aligned} \pi_C(s) &= \left(n^{Out}_{CC}(s) + n^{In}_{CC}(s)\right) - b\left(n^{Out}_{CD}(s) + n^{In}_{CD}(s)\right), \\ \pi_D(s) &= (1+a) \cdot \left(n^{Out}_{DC}(s) + n^{In}_{DC}(s)\right). \end{aligned}$$

Expected normalized lifetime utilities are then simply $u_X = (1-\delta)\sum_{s=0}^{\infty} \delta^s \pi_X(s)$, $X \in \{C, D\}$. Simplifying the expressions and scaling by the factor $1/K$ delivers

$$u_C = \frac{2pq - b(1-q)(1+p-\delta^2(1+p(1-pq)))}{1-\delta^2(1-pq)}, \tag{1}$$

$$u_D = (1+a)\left(pq + q\frac{1-\delta^2}{1-\delta^2(1-pq)}\right). \tag{2}$$

We remark that δ plays two distinct roles in the model. First, it determines the turnover rate at which agents enter and leave the system. Because of this, δ has a direct effect on the evolution of the system, holding fixed the behavior of all agents. It is in this role only that δ appears in our analysis until we come to the computation of u_C and u_D. Second, δ affects the preferences of agents because it represents the effective discount factor. Thus for any given system dynamics, δ influences optimal behavior.

3.3 Characterization of simple stationary equilibria

Each agent chooses at birth C or D so as to maximize his expected utility. In order to characterize optimal choices under consistent strategies, we are interested in comparing u_C and u_D as a function of q and p under various parameterizations of the model. It is convenient to define $V(q, p; a, b, \delta) = u_C - u_D$.

For any given steady-state $L_{(p,q)}$, the expected value to a cooperator of a relationship from a new inlink is proportional to

$$v = \frac{r_C}{1 - \delta^2} - br_D.$$

Using the expressions above, we conclude that v is non-negative if and only if

$$b \leq \left(\frac{q}{1-q}\right)\left(\frac{1}{1 - \delta^2(1 - pq)}\right). \tag{3}$$

The following observations describe the various possibilities for SSE. First, notice that for any choice of parameters, there will exist an SSE with $q = 0$. This is true because when all agents in the system defect, defection is strictly optimal, i.e., $V(0, p; a, b, \delta) < 0$. For some parameters, this will in fact be the unique SSE, in which case it is not possible to sustain any level of cooperation.

Next, notice that at any SSE (p, q) where $q > 0$ (there exist cooperators) it must be that $p > 0$ (cooperators accept some inlinks) so that equation (3) must hold. If the SSE is such that $p < 1$, then equation (3) must hold with equality, leaving cooperators indifferent to the acceptance of inlinks.

If the SSE is such that $q = 1$ (there is universal cooperation), then it must be that $p = 1$ (accepting inlinks is dominant). If the SSE involves an interior solution for q, then it must be that $V = 0$, as entering agents must be indifferent between cooperation and defection. In this case, the SSE might or might not involve an interior solution for p; this depends on the parameters.

We are particularly interested in those SSE that are *stable*.

Definition **3** *A simple stationary equilibrium (p, q) is stable if there exists $\epsilon > 0$ such that*

(i) if $q < 1$, then for all $q' \in (q, q + \epsilon)$, $u_C(p, q') < u_D(p, q')$, and

(ii) if $q > 0$, then for all $q' \in (q - \epsilon, q)$, $u_C(p, q') > u_D(p, q')$.

Our main result of this section establishes uniqueness of stable SSE.[15] The autarky outcome ($q = p = 0$) constitues a stable SSE for all parameters.

THEOREM **1** *If there exist SSE with $q > 0$, then there is in fact a unique such stable SSE.*

Proof. See Appendix. □

The proof is involved but expresses an intuition that we summarize here. The set of points $(q, p) \in [0, 1]^2$ in which $V \geq 0$ is a connected (but not convex) region which we call Γ. The set of (q, p) at which $v = 0$ is defined by a strictly increasing function $p_t(q)$ that is negative for q near zero and greater than one for q near one. Define thus its restriction to the unit square by $t(q) = \min\{\max\{p_t(q), 0\}, 1\}$. SSE occur (i) at the intersection of $t(q)$ with the boundary of Γ and (ii) at $(q, p) = (1, 1)$ when $(1, 1) \in \Gamma$. When $(1, 1)$ is an SSE it is generically stable, but otherwise stability involves the extra requirement that the intersection is on the right boundary of Γ and not the left boundary. The proof proceeds by limiting the number and type of possible intersections between Γ and t through explicit consideration of the utility functions.

We view stability as an important refinement in our setting. SSE that fail the stability requirement are arguably unsatisfactory solutions, and in this sense the model predicts a unique outcome. Nevertheless the structure of the set of all SSE is informative. In the course of proving Theorem 1 we prove the following result.

Proposition 1 *The following statements are true:*

- *All SSE are ordered, in the sense that if (p, q) and (p', q') are two SSE with $q' > q$, then $p' \geq p$, with strict inequality when $p < 1$.*
- *In addition to the autarky outcome, there are generically either 0, 2, or 4 SSE.*
- *If there are 4 additional SSE, the largest one is stable and involves $p = 1$.*

[15] The stability notion is intentionally a weak requirement to emphasize that all other SSE fail even this basic requirement. Stronger notions of stability either agree with our notion or fail existence. Every SSE that is stable under our definition satisfies as well the analogous requirement for p.

- *If there are 2 additional SSE, the smaller one is unstable.*

Proof. See Appendix. □

We outline here the main arguments relevant to Proposition 1. One can show explicitly that the boundary of Γ intersects any given horizontal line of the form $p = \kappa$ at most twice. This means that, in particular, there can be at most two SSE with $p = 1$. In this case, it is clear that the larger of the two SSE is stable, and the smaller is unstable. Through a separate argument using a change of variables we show that $p_t(q)$ can intersect the boundary of Γ at most twice, which in turn limits the number of interior equilibria to at most two. The most difficult part of the proof involves showing that the point in Γ with the smallest value of p lies below $p_t(q)$, which can then be used to show that all but the largest SSE are necessarily unstable.

Note finally that in the case of two interior SSE, the larger one may be stable or unstable, and it may involve $p < 1$ or $p = 1$; if it involves $p = 1$, then it must be stable. Importantly, in all cases when a stable SSE exists, it is the maximal and, therefore best, among all SSE.

The model allows for an explicit determination of the set of SSE and the determination of stability for any parameters (a, b, δ). It is therefore possible in principle to partition the parameter space into regions that map into the different configurations of SSE described by Proposition 1. However, such an analysis is cumbersome. Instead, we present results for the limiting case in which players becoming perfectly patient (i.e., long lived), which captures much, but not all, of the richness of the model. We then exhibit some examples with intermediate values of δ.

Proposition 2 *The following statements hold in the limit as δ approaches one.*
If $a < 1$ then $p = q = 1$ is a stable SSE.
If $a > 1$, then

- *all SSE have $q < 1$*
- *There exists a stable SSE if and only if $b < 1 + a$.*
- *If $(1 + a)/a < b < 1 + a$ the stable SSE has $p < 1$.*

- *If $b < (1+a)/a$ the stable SSE has $p = 1$.*

Proof. See Appendix. □

Proposition 2 formalizes a number of our previous observations. Namely, there is a rich set of qualitatively distinguished outcomes, even in the limit as players become completely patient. Indeed, full cooperation can be achieved for perfectly patient players only if the temptation payoff is $a < 1$.[16] The intuition for this bound comes from considering a world with full cooperation, $q = 1$, and a high discount factor. In this case, a defector earns $1 + a$ per outlink per period. A cooperator earns 1 from each outlink per period, and also expects to build an asymptotically complete neighborhood well before she dies, and in this case has as many inlinks as outlinks, for a total of 2 per outlink per period. Thus if $a > 1$, full cooperation cannot be sustained.

When full cooperation cannot be sustained, it is still possible that partial cooperation can be sustained. In this case society consists of cooperators and defectors co-existing in a specific ratio. This possibility allows for a novel description of some real-world communities such as, perhaps, eBay, in which most transactions are conducted in good faith but where one expects a fringe of cheaters. In our analysis, such an outcome obtains when the temptation payoff is high; patience is not necessarily enough to overcome this effect. If b is small enough that there is a stable SSE where all inlinks are accepted ($p = 1$), then the stable level of cooperation is easily determined to be $\frac{2-b}{1+a-b}$.[17]

Finally, it may be the case that the stable SSE requires exclusivity ($p < 1$). We find this possibility to be of particular interest, since the presence of cooperators is made possible only when those cooperators limit their exposure to society. Clearly this is costly, since some relationships are not materialized. The idea behind exclusivity is the following. If cooperators accepted all proposed links, defection would be relatively attractive, and so to balance the incentives between cooperating and defecting, the level of cooperation would have to be relatively low. But at low levels of cooperation, the expected value of an inlink to a cooperator is negative, and so cooperators would prefer to reject proposed links, leading

[16] This result stands in contrast to much of the work on repeated prisoner's dilemma with random matching, in which the goal is almost always to construct equilibria that always support full cooperation for patient players.

[17] It is worth noting that because of the constraint $a - b < 1$, this possibility obtains only when $a < 1 + \sqrt{2}$.

to a breakdown of the network. The only equilibria therefore involve the rejection of some proposed links. While this reduces the utility to all players, it can reduce the incentive to defect, such that at strong enough levels of exclusivity, a positive level of cooperation becomes consistent with players' incentives.

4 Consistency as self-enforcing behavior

Simple stationary equilibria are defined in a setting that requires agents to apply strategies that are consistent, with the immediate implication that agents are trusting and unforgiving. This can be thought of as an equilibrium that arises under a very natural social norm, in the spirit of [8]. The social norm specifies how to behave in one's relationships as well as how to manage these relationships.

In this section we dispense with the presumption of consistency and study optimal behavior in the absence of social norms that restrict strategies. We find that, under an appropriate parametric condition, the behavior described above is self-enforcing, in the sense that it constitutes play on an equilibrium path.

Before stating the result, we present the condition that will be required. The condition involves p and q as well as (a, b, δ). As it involves the endogenous values of p and q as well as (a, b, δ), the condition should be interpreted as a requirement of a particular steady-state (p, q) under consideration. The condition, which we call the *consistency inequality*, is the following.

Definition 4 *The* consistency inequality *is*

$$\frac{(1+b)(1-pq) - bq(1-p)}{(1+a)(1-pq)} \geq 1 - \delta^2(1-pq). \tag{4}$$

The result we provide is that every SSE (p, q) that satisfies the consistency condition arises as an equilibrium outcome. On the equilibrium path, agents behave in a way that conforms with the norms of being consistent, trusting, and unforgiving. Each agent applies a strategy that, at every history, optimizes its continuation utility in expectation over future randomness and his beliefs about the state of the system given his observations. In turn, these beliefs are consistent with the strategy being employed (recall our focus on symmetric equilibria). A formal development of the solution concept is provided in the Appendix.

THEOREM **2** *Suppose that* $(p, q) \in [0, 1]^2$ *is a simple stationary equilibrium at which the consistency condition is satisfied. Then there is an equilibrium such that if the system is in state* $L_{(p,q)}$*, all agents apply actions that are consistent, trusting, and unforgiving on the equilibrium path. Moreover, under this strategy* q *is a stationary level of cooperation and* p *is a stationary level of inlink acceptance.*

Proof. Section 4.3 is dedicated to the proof. □

4.1 Maintenance of relationships

If agents apply consistent strategies, the beliefs of an individual regarding the future play of his partners are easy to describe. If other agents behave consistently, it is optimal to always maintain a relationship after observing cooperation, and it is optimal to sever a relationship after observing defection. Indeed, these decisions are *strictly* optimal: maintaining a link with a cooperator has positive expected utility, and maintaining a link with a defector has negative expected utility (for a cooperator). A link between two defectors must be severed because the sponsor of that link strictly prefers to re-match and obtain probability pq of interacting with a cooperator at the next period. Importantly, this behavior is sequentially rational and holds for off-path play in which an agent's partner behaves inconsistently, since consistency is defined in terms of taking the same behavior as was taken in the previous period. Under the definition of equilibrium that is detailed in the Appendix, equilibrium beliefs require an agent to assign probability one to consistent behavior of his partners even after observing an off-path inconsistent choice, through use of a standard perfection requirement. It is therefore the case that every best response to a consistent strategy has the property that a link is broken if and only if a defection is observed on that link.

4.2 Consistent Behavior

The analysis in Section 3 was conducted under the assumption that individuals have available to them only two (pure) strategies at their birth governing their choices of cooperation and defection. Optimality, then, requires taking expectations over the implied outcomes of these two actions and choosing appropriately. There is no consideration of deviations from consistency; the choice is assumed to be made with commitment. We now want to show that

if agents play consistent and unforgiving strategies, and the consistency condition is satisfied, then consistent behavior is (part of) a best response. This requires showing that there is no history at which an agent can profitably deviate through the use of an inconsistent action.

First notice that for a defector in a steady-state, the calculation is identical at every round. This is so because, under unforgiving strategies, he loses all of his connections at every period. Thus, if the continuation value of perpetual defection exceeds that of perpetual cooperation at some period, the same conclusion is true at subsequent periods.

For a cooperator the situation is complicated by the fact that the number of relationships with other cooperators changes over time. At an SSE with $q > 0$, a cooperator is at least as happy with his choice, at birth, than he would be under the alternative plan of defection. But, in principle, with positive probability there may arise histories at which a cooperator prefers to deviate by defecting (after which its optimization problem is identical again to the one at birth).

We now introduce notation to describe the state of an individual of age s. For a given agent, let K_s^I denote the number of in-links from cooperators at the beginning of round s, and let K_s^O denote the number of out-links to cooperators at the beginning of round s (i.e., those links that are maintained from the previous period).

The next result provides the key implication of the consistency condition that we will use below to guarantee that cooperators never have a profitable deviation.

Lemma **1** *Suppose $(p, q) \in (0, 1]^2$ is a simple stationary equilibrium and the consistency inequality holds at (p, q). Consider an agent that has K_s^I inlinks and K_s^O outlinks at the beginning of round s, with $K_s^I + K_s^O > 0$. Then the expected utility of cooperating on all rounds starting at s is strictly greater than the expected utility of defecting on round s and then cooperating on all subsequent rounds when other agents play simple strategies.*

Proof. We focus attention on a fixed agent i. Let ϕ_C denote the simple strategy in which agent i cooperates each round and accepts inlinks with probability p, and let ϕ_D denote the simple strategy in which agent i defects each round (and accepts all inlinks). Let ϕ_F denote the strategy in which the agent defects for one round, then cooperates on every subsequent round (and is unforgiving and trusting on every round, accepts all inlinks on the first round, and accepts inlinks with probability p on all subsequent rounds). For an

arbitrary age s and a given strategy ϕ, write $u(\phi, k_I, k_O)$ for the expected utility, evaluated at the beginning of round s, of applying strategy ϕ when $K_s^I = k_I$ and $K_s^O = k_O$, and other players use simple strategies defined by (p, q). To prove the lemma, we must show that $u(\phi_C, k_I, k_O) > u(\phi_F, k_I, k_O)$ whenever $k_I + k_O > 0$.

We will first show that $u(\phi_C, 0, 0) \geq u(\phi_F, 0, 0)$. To see this, note that ϕ_D and ϕ_F are identical on their first round of play, and at the end of that first round agent i will have no links (since other agents apply unforgiving strategies). After that first round, ϕ_F proceeds in the same way as ϕ_C. Moreover, since (p, q) is a simple stationary equilibrium with $q > 0$, we know that $u(\phi_C, 0, 0) \geq u(\phi_D, 0, 0)$. Putting this together, we have

$$u(\phi_F, 0, 0) - u(\phi_D, 0, 0) = \delta(u(\phi_C, 0, 0) - u(\phi_D, 0, 0)) \leq u(\phi_C, 0, 0) - u(\phi_D, 0, 0)$$

from which we conclude $u(\phi_C, 0, 0) \geq u(\phi_F, 0, 0)$.

Write $\Delta u(\phi, k_I, k_O)$ for $u(\phi, k_I, k_O) - u(\phi, 0, 0)$, the utility gain due to adding k_I in-links and k_O out-links to agent i before applying strategy ϕ. We next show that $\Delta u(\phi_C, k_I, k_O) > \Delta u(\phi_F, k_I, k_O)$ for all $k_I + k_O > 0$, which will complete the proof. We note that these utility gains are additively separable in k_I and k_O, so that $\Delta u(\phi_C, k_I, k_O) = \Delta u(\phi_C, k_I, 0) + \Delta u(\phi_C, 0, k_O)$ and $\Delta u(\phi_F, k_I, k_O) = \Delta u(\phi_F, k_I, 0) + \Delta u(\phi_F, 0, k_O)$. We will therefore analyze these gains separately.

Consider first the utility gain due to in-links. We have $\Delta u(\phi_F, k_I, 0) = (1 + a)k_I$, since the agent gains $(1 + a)$ from each link and loses them after his first defection. When applying strategy ϕ_C, the gain is $\Delta u(\phi_C, k_I, 0) = \frac{k_I}{1-\delta^2}$. This is so because the cooperator gets extra utility for each period of the life of the relationship. We have that $\Delta u(\phi_C, k_I, 0) > \Delta u(\phi_F, k_I, 0)$ whenever $\frac{1}{1-\delta^2} > 1 + a$, which is necessary to sustain cooperation in a simple stationary equilibrium anyway.

We turn now to out-links, where a fraction k_O of the agent's out-links are already matched to cooperators, and the remaining out-links will be matched to the population at random. For strategy ϕ_F, $\Delta u(\phi_F, 0, k_O) = (1 + a)(1 - pq)k_O$. To see this, note that the increase in the number of out-links to cooperators is $k_O + (1 - k_O)pq - pq = (1 - pq)k_O$, and this gain is realized for exactly one period.

For cooperators, $\Delta u(\phi_C, 0, k_O) = \frac{(1-pq+(1-q)b)k_O}{1-(1-pq)\delta^2}$. To see this, consider the expected loss experienced by a cooperator who does not have a link pre-formed to a cooperator. If he is

unable to form a new link to a cooperator, then he suffers a loss of 1 (relative to a cooperator who would derive a utility of 1 from a pre-existing link to another cooperator). Moreover, if he forms a link to a defector instead, he loses an additional b due to the interaction with the defector. The total loss is therefore $(1 - pq) + (1 - q)b$, per link. Finally, for a given outlink this gain is maintained as long as the node survives (probability δ), its cooperating partner survives (probability δ), and the outlink of the node in the scenario without the initial k_O cooperate outlinks is not to a cooperator (probability $(1 - pq)$). These events happen independently and hence have a total probability of $\delta^2(1 - pq)$ yielding the above formula. Thus $\Delta u(\phi_C, 0, k_O) > \Delta u(\phi_F, 0, k_O)$ precisely when the consistency inequality holds, completing the proof. □

By virtue of Lemma 1, under the consistency condition, as an agent accumulates relationships with cooperators, the marginal gain from those relationships is maximized by long-term cooperation, and not by defecting. Thus, if it is optimal to cooperate at birth it is necessarily optimal to cooperate at any future point in its lifetime.

4.3 SSE as equilibrium outcomes: Proof of Theorem 2

We shall construct a symmetric equilibrium with the required properties. Recall that a formal definition of equilibrium appears in Appendix A but, informally, what we require is a strategy ϕ^* and a system of beliefs β^* about the state of the network, such that ϕ^* maximizes expected utility at all histories given beliefs β^*, and β^* is consistent with observations under the assumption that other players apply strategy ϕ^*.

The strategy ϕ^* is as follows. First, if on any round an agent observes a fraction of cooperation other than q, or an aggregate inlink acceptance other than p, the agent accepts all proposed links, defects that round, and breaks all links with observed defectors at the end of the round. Note that this behavior is optimal given that (p, q) is publicly observed and other agents also play according to ϕ^*, since these behaviors are optimal given the belief that all other agents will defect. Otherwise, if the agent observes (p, q), the agent takes consistent, trusting, and unforgiving actions defined by (p, q). At birth, upon observing the state (p, q), he cooperates with probability q, and in that case accepts inlinks independently with probability p; otherwise he chooses to defect (and accept all inlinks).

The associated belief system β^* is straightforward. At birth, the agent believes that the system begins in state $L_{(p,q)}$. The agent continues to believe that the system is in steady-state $L_{(p,q)}$ as long as he observes (p, q) at the end of each period. If a fraction of cooperation other than q is observed, or an aggregate inlinnk acceptance other than p, the agent believes that every other agent will defect on subsequent rounds, since (p, q) is public. Even though we have not provided a full description of an agent's belief about the state of the network, the properties discussed are sufficient to determine whether or not ϕ^* is an optimal strategy.

Under ϕ^* agents are consistent, unforgiving, and trusting provided the system remains in state $L_{(p,q)}$, so that (p, q) is indeed stationary under ϕ^*. It remains to show that applying strategy ϕ^* is optimal given the observation of (p, q) and the belief that other agents play according to ϕ^*. Note first that, under the belief that the system begins in $L_{(p,q)}$ and other agents use ϕ^*, it is rational to believe that the system remains in $L_{(p,q)}$ as long as agents observe (p, q). It is therefore sufficient to demonstrate that ϕ^* is optimal under the belief that the state of the system is described by $L_{(p,q)}$ at all times.

We focus attention on a particular agent i. Write $u(\phi)$ for the expected lifetime utility of agent i when applying strategy ϕ. Let ϕ_{opt} denote a strategy that maximizes expected utility against the profile of all agents playing ϕ^* in state $L_{(p,q)}$, and suppose for a contradiction that $u(\phi_{opt}) > u(\phi^*)$. Let ϕ_C denote the trusting, unforgiving and consistent strategy in which the agent chooses cooperation at birth, and accepts each incoming link independently with probability p, and let ϕ_D denote the similar strategy in which the agent chooses defection and accepts each incoming link.

Note first that if $q = 0$, then no strategy obtains positive expected utility; thus strategy $\phi^* = \phi_D$ is optimal, since in this case $u(\phi_D) = 0$. We therefore assume $q > 0$ for the remainder of the proof.

As discussed in Section 4.1, we know that every optimal strategy breaks a link if and only if a defection is observed on that link. In particular, ϕ_{opt} must satisfy this property. Moreover, every optimal strategy accepts all inlinks on a round in which it prescribes defection so, in particular, ϕ_{opt} must satisfy this property.

For all $r \geq 1$, define the random variable T_r as the age at which ϕ_{opt} prescribes that agent i defect for the r'th time. We then define strategy ϕ_D^r as the strategy in which agent i follows ϕ_{opt} up to and including round T_r, after which point he behaves according to ϕ_C. We

also define strategy ϕ_C^r as the strategy in which agent i follows ϕ_{opt} up until round T_r, but on round T_r and all subsequent rounds he behaves according to ϕ_C. Thus ϕ_C^r and ϕ_D^r differ only on their actions on round T_r, in which ϕ_C^r specifies cooperation (and accepting inlinks with probability p) and ϕ_D^r specifies defection (and accepting all inlinks). For notational convenience we define $\phi_D^0 = \phi_C^0 = \phi_C$.

We first claim that $u(\phi_C^r) \geq u(\phi_D^r)$ for all $r \geq 1$. Strategies ϕ_C^r and ϕ_D^r are identical until round T_r, at which point ϕ_C^r proceeds to cooperate on every subsequent round, whereas ϕ_D^r defects for a single round and then cooperates thereafter. Therefore, Lemma 1 directly implies that $u(\phi_C^r) \geq u(\phi_D^r)$, as agent i maximizes utility by cooperating on round T_r regardless of the number of maintained links with cooperators on round T_r.

We next claim that $u(\phi_D^{r-1}) \geq u(\phi_C^r)$ for all $r \geq 1$. Strategies ϕ_D^{r-1} and ϕ_C^r are identical through round T_{r-1}, after which both strategies prescribe cooperation on each turn, but ϕ_C^r does not necessarily accept every proposed inlink independently with probability p from round $T_{r-1} + 1$ to round T_r. However, by virtue of the assumption that (p, q) is an SSE, it must be optimal to accept proposed links independently with probability p when perpetually cooperating, and thus $u(\phi_D^{r-1}) \geq u(\phi_C^r)$.

Combining these two claims, we have that $u(\phi_D^{r-1}) \geq u(\phi_D^r)$ for all $r \geq 1$. But $\phi_D^0 = \phi_C$, and $\lim_{r\to\infty} u(\phi_D^r) = u(\phi_{opt})$ (noting that the limit must exist since utilities are time-discounted). We therefore conclude $u(\phi^*) \geq u(\phi_C) \geq u(\phi_{opt})$, which is the desired contradiction.

4.4 Application of the consistency condition to SSE outcomes

We now discuss the consistency condition in more detail. Of course, direct application of the consistency condition will verify whether or not any given set of parameters (a, b, δ) and a given steady-state (p, q) permits consistency of cooperation in equilibrium. But the condition is complex enough that it is not transparent to immediately assess how stringent the condition is. The following result quantifies our assertion that the optimality of consistent cooperation is not a dramatically more demanding requirement then the optimality of any cooperation at all.

Proposition 3 *The following statements are each (individually) sufficient for the consistency condition to hold.*

- $p = q = 1$
- $b \geq a$ *and* $p = 1$
- $b \geq a$, (p, q) *is a stable SSE for* (a, b, δ), *and* $\delta \to 1$

Proof. The first two items are easily verified by direct inspection of the consistency condition. The last item is proved in several cases. Consider the stable SSE described in Proposition 2 which characterize the limit case of interest, $\delta \to 1$. If $a < 1$ then the stable SSE is $(1, 1)$ and covered by the first case. So assume $a > 1$. By Psroposition 2 there exists a stable SSE only if $b < 1 + a$. If $(1 + a)/a < b < 1 + a$ then the limiting stable SSE is interior, and one can easily solve for it explicitly to obtain $(p, q) = (\frac{1+a+b}{(1+a)b}, \frac{b}{1+a-b})$. It is obvious that the consistency condition becomes tighter as b decreases, and so it is enough to verify that it holds at the presumed lower bound of $b = a$, in which case (p, q) reduces to $(\frac{1+2a}{a(1+a)}, a)$. Making these substitutions for (p, q) into the consistency condition produces an inequality that is trivially satisfied. Finally, when $b < (1 + a)/a$ the limiting stable SSE has $p = 1$ and $q = \frac{2-b}{1+a-b}$. Thus it is covered by the second item, but is also easy to verify directly. Substituting into the consistency condition and evaluating as $\delta \to 1$ yields $b \geq a + \frac{1-\sqrt{4a^2-3}}{2}$, which is implied by $b \geq a$ whenever $a \geq 1$, as desired. □

Notice in particular that if the stage game is supermodular ($b \geq a$) then at least in the limiting case as players become perfectly patient, all stable SSE satisfy the consistency condition. In particular, this verifies that under supermodularity, all of the stable SSE described by Proposition 2 survive as equilibria in which consistent, trusting, and unforgiving behavior is self-enforcing.

5 Inconsistent behavior

We have argued that when the consistency condition is satisfied, SSE outcomes are supported as on path behavior of an equilibrium without any presumption of consistency (Theorem 2). Further, we have argued that the consistency condition is not an overly stringent requirement

(Proposition 3). We now take up the task of analyzing behavior when the consistency condition fails. The main result of this section asserts that if players are patient and if the network is sufficiently dense, then for every SSE there exists an ϵ-equilibrium with approximately the same levels of cooperation and inlink acceptances.

Towards this end, notice that an agent faces a potential tension in his incentives across different links. In particular, the incentive to defect is stronger along outlinks than inlinks, by virtue of the fact that following a defection and the loss of all relationships, the outlink can be immediately rematched (albeit with probablity only pq of connecting to another cooperator), whereas the inlink is lost permanently with no chance of replacement. Lemma 1 shows that the consistency inequality implies that an agent gains more utility from an outlink to a cooperator by cooperating than by defecting and forming a new relationship, thereby guaranteeing that incentives are aligned towards consistent cooperation. But if the consistency condition is violated, this no longer holds, and exploiting outlinks to cooperators would be profitable.

We conclude that the incentive to defect is strongest when the number of outlinks to cooperators is high relative to the number of inlinks from cooperators. We formalize this conclusion as follows.

Proposition 4 *Consider an SSE (p, q) with $q > 0$. A cooperator has a profitable inconsistent deviation if and only if*

$$\kappa_O \left[(1+a)(1-pq) - \frac{1-pq+(1-q)b}{1-\delta^2(1-pq)} \right] > \kappa_I \left[\frac{1}{1-\delta^2} - (1+a) \right]. \tag{5}$$

Proof. From the expressions in the proof of Lemma 1 it is easily seen that (5) captures the set of pairs (κ_O, κ_I) at which the incentive to defect is stronger than at birth, when $(\kappa_O, \kappa_I) = (0, 0)$. There are two cases to consider. First consider the SSE that has $(p, q) = (1, 1)$, in which case the right hand side of (5) is non-negative (this is true for any SSE with $q > 0$) while the left hand side is zero. Thus (5) never holds in this case and by Proposition 3 a cooperator never has a profitable deviation. Second, consider an SSE with $q < 1$ in which case $u_C = u_D$ at birth. This indifference means that (5) captures exactly the situations in which defection becomes strictly profitable, completing the proof. □

Note that Proposition 4 shows that the consistency condition is tight, in the sense that at an SSE where the consistency condition fails, a cooperator reaches a history at which

he has a profitable deviation with positive probability. To see this, note that with positive probability $\kappa_I = 0$ and $\kappa_O > 0$, in which case the failure of the consistency condition means that the left hand side of (5) is strictly positive.

Consider now an SSE (p, q) at which the consistency condition is violated, and where all agents play the simple strategy (p, q). By Proposition 4, a best response is the following "threshold" strategy: behave exactly like a simple cooperator, with the exception that whenever κ_O/κ_I exceeds the threshold defined by (5) the agent should defect for one round and then restart the strategy exactly as if he was a newborn agent. Now, as δ becomes large, the right hand side of (5) increases without bound, so that for high enough δ the only circumstance under which a cooperate would defect is when $\kappa_I = 0$. Due to a law of large numbers, as K, the number of outlinks per agent, grows the probability of $\kappa_I = 0$ vanishes exponentially fast.

Combining these observations, for patient players and dense networks we can construct an ϵ-equilibrium with actions and payoffs arbitrarily close to those of a given SSE. We formalize this conclusion as follows.

Proposition 5 *Fix (a, b) such that there exists a stable SSE for sufficiently large δ. For any $\epsilon > 0$, there exist a $\bar{\delta}$ and a $\bar{K}$ such that for all $\delta > \bar{\delta}$ and for all $K > \bar{K}$, there exists an ϵ-equilibrium with a stationary level of (p', q') that is within ϵ of the associated stable SSE, (p, q), i.e., $|p - p'| < \epsilon$ and $|q - q'| < \epsilon$.*

Proof. For a given δ let (p, q) denote the stable SSE. The strategy we use is the following. Every agent mixes at birth and plays cooperate with probability q and defect with probability $1 - q$. In the latter case the agent plays exactly like a consistent defector. In the former case the agent accepts inlinks with probability p while cooperating and plays the threshold strategy defined by (5) relative to (p, q). Set $\bar{\delta}$ so that if $\delta > \bar{\delta}$ the only solution to (5) is $\kappa_I = 0$. As K becomes large, the probability that $\kappa_I = 0$ in steady-state vanishes exponentially, so that $\bar{K}$ can be chosen to make the probability of a cooperator deviating at any age as small as desired.

This directly implies several facts. First, because deviations from simple strategies are rare, the implied steady-state is near (p, q). Second, because the threshold strategy differs only rarely from the simple cooperating strategy, its expected utility is very near that of

the simple cooperator at the steady-state, which is in turn very near the expected utility of a simple cooperator at (p, q). Third, the expected utility of a defector is near that of a defector in a world where everyone plays simple (p, q) strategies, because of the fact that the steady-state is very near (p, q) and because the inconsistent behavior of cooperators has no effect on a defector beyond its effect on steady-state level of cooperation.

Therefore, for sufficiently large δ and K, agents applying the proposed strategy, conditional on whether they become defectors or threshold-cooperators, have expected utilities that are arbitrarily near the expected utilities of players using simple strategies at the (p, q) SSE. And because of the observation that the threshold strategy defined by (5) at the (p, q) steady-state is in fact a best response, this strategy has utility as close as desired to an optimal strategy. □

This result establishes that, even in the case where the consistency condition fails, our description of SSE survives as a reasonable description of equilibrium behavior.

6 Conclusion

We have developed a model of interactions for a large anonymous community with turnover, in which agents are interconnected via an endogenously evolving network. The class of simple strategies that involve consistency of choices over time provides the foundation for our analysis. With consistent behavior, the social norm of ostracism is required in equilibrium, whereby links to an agent who defects are always severed immediately. We view this form of ostracism as capturing an empirically relevant phenomenon that is used to support cooperative behavior.

Under consistent strategies, we fully characterize stationary equilibria. Universal cooperation is sustainable for a non-trivial range of parameters, but not always. In particular, it requires not only that players are sufficiently long-lived (i.e., patient), but also that the temptation payoff for defecting not be too large. When these conditions are not both met, the presence of some level of non-cooperative behavior in a large anonymous system is unavoidable.

We believe this captures an important feature of a number of applications for which a fringe of exploitative behavior is observed. Our analysis offers new insights into thinking

about how much cooperation can be sustained as a function of the underlying parameters of the system. When some level of defecetion persists, we can address through comparative statics what kinds of policies could be expected to improve the level of cooperation, and the total welfare of the system. Nearly all other related work falls either into the category of constructing equilibria for which full cooperation is sustainable, or else analyses models in which some form of inefficiency is inevtiable. In a sense, our framework allows for a more balanced description of the acheivable level of cooperation in a society. This characterization is specifically due to the novel tradeoff in our model: that between immediate gains to defection and the gain to accumulating social capital in the form of additional links with cooperators over one's life.

When full cooperation does not obtain, the presence of defectors causes relationships among cooperators to be viewed as a scarce and valuable resource, which we identify as a form of social capital. In this case, the model provides for the possibility of a form of exclusivity to arise endogenously, in which cooperative players only occasionally agree to form new relationships with strangers. This exclusivity comes at a net welfare loss to society, but is necessary to incentivize cooperative behavior, as it slows down the rate at which cooperating partners can be found, thereby strengthening the penalty of ostracism due to defecting.

As it turns out, the characterization of equilibria under consistent strategies says a lot about equilibrium outcomes in which the possibility of inconsistent behavior is allowed. This demonstrates that the simple behavior we focus on can be self-enforcing, a finding that is not obvious a priori. In particular, under the appropriate condition, every simple stationery equilibrium has a corresponding equilibrium that supports consistent behavior on the equilibrium path with the same steady-state level of cooperation. The condition, which is satisfied for many parameter values, requires that the returns to links with cooperators are higher for cooperators than for defectors, thus implying that persistent cooperation is sequentially rational. Finally, we show that when this consistency condition is violated, a "rob the bank" strategy, in which a cooperator deviates when he has a sufficiently low level of social capital, forms an epsilon equiibrium when players are patient and can maintain many links. We view these results as establishing a robust class of equilibria in a co-evolving network that generates a high degree of cooperation.

We have made some of the modeling choices to emphasize the point that it is possible to maintain cooperation even under unfavorable conditions. In thise sense, our results could be viewed as identifying a lower bound on what might be expected to obtain in related models. For example, one could imagine that new partners are found both at random, as we have modeled, and by searching the neighbors of current partners, as in [10]. This would allow cooperators to preferentially find other cooperators more quickly, and would tilt incentives in favor of cooperation. It would also bring the degree distribution closer in line with the empirical observation that social networks tend to exhibit heavy tails. One could also imagine that agents have less than perfect access to anonymity. This generally has the consequence of making punishments for defection stronger, thereby increasing the scope for cooperation. Similarly, it might be reasonable to assume that cooperators have longer average lives, which again makes cooperation easier to sustain.

We conclude with a remark about welfare. While we have not provided a complete characterization of equilibria in our framework, there is reason to be optimistic that the simple consistent equilibria that we identify generate a high level of average utility relative to other potential equilibria. The reason is that any strategy that incentivizes cooperation through inconsistent strategies has the difficulty that defecting with one partner requires defecting on all partners simultaneously. As such, the incentives of how to behave in the context different relationships are potentially in conflict, resulting in either diminshed incentives for cooperation or the inefficient loss of relationships. Consistent strategies have the unique property that an agent is never called on to change his behavior over the course of a relationship. We conjecture that the welfare associated with the stable SSE is maximal among all equilibria.

A Appendix A: formal development of the model

The model described in this paper is somewhat complex, incorporating a changing set of players, a large state space that is almost entirely unobserved by each individual player, and various sources of randomness. In the main text, we approached this model by handling the notions of strategies, equilibria, and beliefs in an informal manner. In this appendix we redescribe these concepts more formally, which will allow us to state the results more precisely.

A.1 Histories and Actions

The strategy of an agent is a mapping from its (private) history to (a probability distribution over) actions. The history encodes all the information the node has acquired during its life. In particular, the history of an agent contains its observation of (p, q) at each point during its life, all of its past actions, and the actions of each of its partners over time, together with how and when those relationships were initiated and ended. The action space at any history is a choice of C or D, together with whether or not to sever any existing relationships and accept any new proposed links. We now develop these elements more formally.

The set of agents is the unit interval $N = [0, 1]$. Whenever an agent dies, it is replaced by an agent who takes the same name. We focus on an arbitrary agent i. Denote the age of i by s. In the period when i is born, $s = 0$; s increments by one in each subsequent round in which i remains alive. At each point in time, i observes the value of (p, q) determined by the choices at the previous round. Define (p^s, q^s) to be the proportion of cooperators and rate of inlink acceptance that i observes in the round when i is age s. At each s, i chooses an $\alpha_i \in \{C, D\}$. For each partner j that i has at age s, the vector $\beta_j^s = \{\alpha_j^s, d_j^s, e_{ji}^s, e_{ij}^s\}$ defines the action that j takes, and whether and, if so, how the link was terminated in that round. The variable d_j^s equals 1 if j dies (0 otherwise), and the variables e_{ji}^s and e_{ij}^s record whether j or i respectively chooses to sever the link (a value of 1 corresponds to severing, 0 to not severing).

The collection of i's partners is recorded in two lists. The outlinks of agent i at age s are stored in a vector Out_i^s of length K. If i's k'th outlink at age s is to agent j, then the k'th element of this array is β_j^s. Due to anonymity, though, i does not know the value of j, but only the values of the elements in β_j^s. The inlinks of agent i require a bit more notation since there is not a fixed number of them. To account for this, we define a vector In_i^s representing the state of all current and past inlinks of agent i at age s. The k'th component of this list records information pertaining to the k'th inlink proposed to agent i over his life. Initially, In_i^s is empty. When agent i at age s receives a proposal for an inlink from agent j, it updates In_i^s as follows: if the link is accepted it appends β_j^s to In_i^s; if the proposed inlink is rejected outright, then we append a special symbol REJECT to In_i^s. After actions are realized, agent i updates each β_j^s in In_i^s appropriately. We define the *size* of list In_i^s, denoted

by $|\text{In}_i^s|$ to be the number of *active* links contained in the list, i.e., the number of components of In_i^s for which $d_j^s = e_{ji}^s = e_{ij}^s = 0$.[18] Again, it is important that i not know the values of j corresponding to the various inlinks in In_i^s. Finally, denote by L_i^s the number of inlinks proposed to i in round s.

The information that i collects from the round in which he is age s is

$$h_i^s = \{p^s, q^s, \alpha_i^s, L_i^s, \text{Out}_i^s, \text{In}_i^s\}.$$

The (private) history of i at age s is the vector $H_i^s = \{h_i^0, \ldots, h_i^s\}$. In a valid history it must be the case that the length of the list In_i^s grows monotonically with s and that if the k'th component of In_i^s is either REJECT or a β_j^s indicating a link termination (i.e., either d_j^s, e_{ji}^s, or e_{ij}^s equals 1), then this component remains constant for the remainder of i's lifetime (i.e., for all $t > s$, the k'th component of In_i^t equals the k'th component of In_i^s). Denote the space of feasible age-s histories for i by $\mathcal{H}_i^s$. The set of all histories for i is then $\mathcal{H}_i = \cup_s \mathcal{H}_i^s$.

At each round, i takes three separate actions: (i) the choice of α_i, (ii) the acceptance or rejection of proposed inlinks, and (iii) the severance or continuation of each active link. The (history dependent) action set of i at age s is $A_i^s(H_i^s) = [0,1] \times [0,1]^{L_i^s} \times [0,1]^{K+|\text{In}_i^s|}$, with the interpretation that the first element specifies the probability that i chooses C at age s, the second element specifies the probability of accepting each proposed inlink, and the final element specifies the probability that i severs a link to each of his partners.

Let $\mathcal{A}_i^s = \cup_{H_i^s \in \mathcal{H}_i^s} A_i^s(H_i^s)$ denote the set of all age-s action sets, and let $\mathcal{A}_i$ denote the space of all action sets for i.

A strategy for i is a mapping $\phi_i : \mathcal{H}_i \to \mathcal{A}_i$, with the restriction that $\phi_i(H_i^s) \in A_i^s(H_i^s)$ for all $H_i^s \in \mathcal{H}_i$. When i makes the choice of α_i^s, he has all the information in H_i^{s-1} as well as (p^s, q^s), but he has not observed the remainder of h_i^s. Similarly, when i makes his choice of accepting inlinks, he observes h_i^{s-1} and $(p^s, q^s, \alpha_i^s, L_i^s)$, but nothing else from round s. Last, when i makes the choice of severing active links, he has observed, additionally, the actions $\{\alpha_j^s\}$ in round s of each of his active partners. We place the associated restrictions on strategies, so that actions depend only on the information observed at each of these times within a round.

[18] Note that one can analogously define the size of Out_i^s; however as agents always replace outlink partners instantaneously, $|\text{Out}_i^s| = K$ for all i and s.

Notice that, implicit in the construction of strategies is the Markovian property that, while actions generally depend on the age of an agent, they cannot be conditioned explicitly on time.

A.2 Equilibria

Recall that, in our definition of histories and actions, a single round involves a sequence of action choices to be resolved by an agent, where incremental observations are made between each choice. We then ensured that a strategy can use only the "currently available" information from the latest round of a history when defining action choices. While consistent with our informal game description, this point of view is notationally cumbersome. A change of variables would allow us to consider each step of a round as a separate information set, in which case a strategy is a mapping from histories to actions without restrictions. We will proceed with our discussion under this change, with the understanding that our notion of a history, strategy, etc. are fully equivalent to those developed in the previous section. Notice that $\mathcal{H}_i = \mathcal{H}^*$ and $\mathcal{A}_i = \mathcal{A}^*$ for all $i \in N$.

A *state of the world* ω is a directed graph with (labeled) vertex set $N = [0, 1]$, plus a history for each vertex. A state represents the links between players in a given round, along with each of their past observations. We write Ω for the set of all possible states of the world. In general, given any set S, we will write $\Delta(S)$ for the set of probability distributions over S.

A *belief* for agent i is a function $\beta_i : \mathcal{H}^* \to \Delta(\Omega)$ that maps each observed history to a distribution over possible world states. We interpret $\beta_i(H_i)$ as capturing agent i's beliefs about the state of the world given a sequence of observations.

We focus on strategy and belief profiles that are symmetric across agents, i.e., there is some strategy ϕ and belief β such that $\phi_i = \phi$ and $\beta_i = \beta$ for all $i \in N$.

Our goal is to define a notion of a symmetric equilibrium, which will be a pair (ϕ, β) that satisfies certain properties. Informally, we wish for the following: at all valid histories ϕ maximizes expected utility given β when other agents apply ϕ; β is consistent with an agent's observations and with the belief that all agents apply strategy ϕ; and, when faced with an unexpected history, β maps to a limit point of beliefs under a vanishing tremble probability. We now describe each of these desiderata in more detail.

We write $u_i(\bar{h}_i)$ for the expected continuation utility obtained by agent i, where $\bar{h}_i$ denotes a distribution over future histories that i will observe. Note that $\bar{h}_i$ captures any dependency on the strategy employed by agent i, as it is a distribution over future observations. Given strategies ϕ, ϕ_i' and state ω, we write $h_i^{\phi}(\phi_i', \omega) \in \Delta(\mathcal{H}^*)$ for the distribution over all future histories that will be observed by agent i when agent i applies strategy ϕ_i' and all other agents apply strategy ϕ, starting from state ω. We extend h_i^{ϕ} to accept a distribution over states in the natural way. We then say that ϕ is *optimal under belief* β if, for all $H_i \in \mathcal{H}^*$,

$$\phi \in \arg\max_{\phi_i'} \{u_i(h_i^{\phi}(\phi_i', \beta(H_i)))\}.$$

That is, for every history H_i, ϕ maximizes the expected utility of agent i given the distribution $\beta(H_i)$ over states, under the assumption that other agents apply strategy ϕ. We also say that ϕ is δ-approximately optimal if for all $H_i \in \mathcal{H}^*$, $u_i(h_i^{\phi}(\phi, \beta(H_i))) \geq u_i(h_i^{\phi}(\phi_i', \beta(H_i))) - \delta$ for all alternative strategies ϕ_i'.

Given ϕ, we now define the *progression function* $P^{\phi} : \Delta(\Omega) \to \Delta(\Omega)$. Given $\sigma \in \Delta(\Omega)$, $P^{\phi}(\sigma)$ is the distribution over states that results when all agents apply strategy ϕ for one round, starting from a state drawn from σ. Note that the resulting distribution is taken over randomness in strategy ϕ and the randomness inherent in the model, i.e. the death and matching processes. We next add the effects of an agent's observations to this distribution: given a distribution $\sigma \in \Delta(\Omega)$ over states, an agent i, and an observation h_i from a single round, we define $P^{\phi}(\sigma, h_i)$ to be the distribution over states that results after resolving a single round of play under ϕ, starting at a state drawn from σ, given that agent i observes h_i in that round. Note that this distribution is well-defined: one can consider the probability of observing h_i given each possible state and apply Bayes' rule.

We say that β is *consistent with strategy* ϕ if, for all i, s, H_i^{s-1} and h_i^s,

$$\beta(H_i^s) = P^{\phi}(\beta(H_i^{s-1}), h_i^s).$$

Observe that the requirement that β be consistent with strategy ϕ does not impose any restrictions on beliefs upon observation of a history that is inconsistent with ϕ. Thus, if this condition is taken to be sufficient for characterizing permissible equilibrium of beliefs, we have the undesirable feature that beliefs and, hence, behavior, is not appropriately restricted off the equilibrium path. This motivates us to require a form of perfection. Given an unexpected

history H_i that has zero probability under ϕ, we would like (informally speaking) for agents to place belief in a minimal number of deviations from ϕ that yield a state consistent with H_i. To achieve this property formally, we will require not only that β be consistent with the application of strategy ϕ by all agents, but also that it maps to a limit point of beliefs under a vanishing trembling probability on actions.

We now formalize the intuition described above. Given any strategy ϕ and any $\epsilon \geq 0$, the ϵ-perturbation of ϕ is the strategy ϕ^ϵ that, independently for each action, follows ϕ with probability $1-\epsilon$, and with the remaining probability chooses an action uniformly at random. We say that β is robustly consistent with ϕ if

- β is consistent with ϕ,
- for all $\epsilon > 0$, there exists belief β^ϵ such that β^ϵ is consistent with ϕ^ϵ, and
- $\lim_{\epsilon \to 0} ||\beta^\epsilon - \beta||_{TV} = 0$ where $|| \cdot ||_{TV}$ denotes total variation distance.

Note that if ϕ is optimal given β, and β is robustly consistent with ϕ, then (taking β^ϵ as in the definition of robust consistency) ϕ^ϵ must be δ-approximately optimal for β^ϵ, where $\delta \to 0$ as $\epsilon \to 0$.

We are now ready to define our equilibrium concept. We say that (ϕ, β) is an equilibrium if ϕ is optimal given β, and β is robustly consistent with ϕ. Note that such an equilibrium always exists. For example, the ϕ that maps every history to "always defect" (formally, using the notation from the previous section, for all $H_i^s \in \mathcal{H}_i^s, \phi(H_i^s) = 0 \times [1]^{L_i^s} \times [1]^{K+|\mathrm{In}_i^s|}$), is a trivial equilibrium.

B Appendix B: Proofs of results for Section 3

We drive towards proving uniqueness of stable SSE with $q > 0$. The auxiliary results of Propositions 1 and 2 are proven in the course, and duly noted. We use $d = \delta^2$ for convenience.

B.1 Acceptance of inlinks

For a cooperator, indifference between accepting and rejecting a given inlink requires that $\frac{r_C}{1-d} - br_D = 0$. Solving this condition for p yields

$$p_t(q) = \frac{q - b(1-q)(1-d)}{bq(1-q)d}.$$

By inspection, this function is negative for small q (it approaches $-\infty$ as $q \to 0$), greater than one for large q (it approaches ∞ as $q \to 1$), and strictly increasing in q on the unit interval. Define, thus,

$$t(q) = \min\{\max\{0, p_t(q)\}, 1\}$$

as the constrained solution to the threshold for accepting inlinks. $t(q)$ is increasing. Note that steady-states (p, q) for which $p < t(q)$ have the property that accepting inlinks is dominant for cooperators. SSE (p, q) require that $p = t(q)$.

This observation implies the first item of Proposition 1.

B.2 Cooperation and Defection

Indifference between cooperation and defection requires that $u_C = u_D$.

Define $V = u_C - u_D$ and $\Gamma = \{(q, p) \in [0, 1]^2 | V \geq 0\}$ as the set of steady states in which cooperation is optimal.

We claim that $V =$ is strictly increasing in d for all $(p, q) \in [0, 1]^2$. To see this, differentiate V with respect to d to obtain

$$(q(1+a) + b(1-q) + 2(1-pq)) \frac{pq}{(1-d(1-pq))^2}, \tag{6}$$

which is strictly positive for $(p, q) \in [0, 1]^2$. This implies that Γ is strictly increasing (in the sense of set inclusion) in d.

To describe the boundary of Γ, solve $V = 0$ for p to obtain:

$$p = \frac{A \pm \sqrt{B}}{C},$$

where

$$\begin{aligned}
A &= q(1+d) - [aq + b(1-q)](1-d), \\
B &= [b(1-q)(1-d) + q(a-1-(a+1)d)]^2 - 4q(b + q(1+a-b))^2 d(1-d), \\
C &= 2q(b + q(1+a-b))d.
\end{aligned}$$

Let us call these solutions $p_1(q)$ and $p_2(q)$ so that we have $p_1(q) \geq p_2(q)$ in the unit square when the solutions are real.

The solutions are real when $B \geq 0$. Notice that B is cubic in q with a leading coefficient of $-4(1+a-b)^2 d(1-d) < 0$. Thus, ignoring the constraint that $0 \leq q \leq 1$, $B \geq 0$ for sufficiently small q and possibly also for a finite interval of q. We have that $B(q=0) = b^2(1-d)^2 > 0$ and that $B'(q=0) = 2b(1-d)[a(1-d) - (1+b)(1+d)] < 0$ whenever $a - b < 1$, as we assume. But for q near zero, the solutions to p_1 and p_2 are not in the unit interval, which can be verified directly, and so these are not valid solutions to $u_C = u_D$; in this case $u_D > u_C$.

In fact, it is the values of q that lie between the second and third roots of B that define the boundary of Γ. To see this, note that as $d \to 1$, $p_1(q) \to \frac{2}{b+q(1+a-b)}$ and $p_2(q) \to 0$. Thus as $d \to 1$, $\Gamma \to \{(q,p) \in [0,1]^2 | p \leq \frac{2}{b+q(1+a-b)}\}$, and it converges in a way that the leftmost point approaches $q = 0$ from the right, while the first root of B is strictly between 0 and the second root. Thus, for any (q,p) with small q, there is a d such that q is equal to the second root of d, and for slightly smaller d, q falls between the two roots, so that $u_D > u_C$ for that value of q and d, independent of p. Then, because Γ is increasing in d, it must be that Γ excludes all such values of q for all smaller d. So valid solutions occur between the second and third roots of B, after taking the intersection with the unit interval.

Ignoring the constraint to the unit square in (q,p) space, we then see that $p_1(q)$ and $p_2(q)$ form the boundary of a connected region. Points inside that region correspond to (q,p) for which $u_C \geq u_D$.

B.3 Number of equilibria

We already described Γ as $d \to 1$. Notice that this directly implies the first claim in Proposition 2.

Consider the system of equations $V = 0$ and $p = \kappa$, for an arbitrary constant κ. It is easy to explicitly solve these equations and see that they have at most two solutions in q; they are of the form $q = \frac{A' \pm \sqrt{B'}}{C'}$. This means that the boundary of Γ intersects any given horizontal line at most twice. In particular, this implies that there are at most two equilibria involving $p = 1$. It also implies that $p_1(q)$ and $p_2(q)$ are single-peaked on the unit square.

We now show that, in addition to the $p = 1$ equilibria, there exist at most two interior equilibria. To accomplish this we use a change of variables from (q,p) to (x,y) where $x = q$

and $y = pq$. Interior equilibria must satisfy both $V = 0$ and $p = p_t(q)$. This system can be written as

$$\frac{2y - (1-d)f(x)}{g(y)} - f(x)y/x = 0 \tag{7}$$

$$\frac{x}{g(y)} - b(1-x) = 0, \tag{8}$$

where $f(z) = b + (1 + a - b)z$ and $g(z) = 1 - d(1 - z)$.

Substituting (8) into (7) and simplifying produces

$$2y - (1-d)f(x) - (1+a)yg(y) - y = 0, \tag{9}$$

Interior equilibria are thus described by simultaneous solutions to (8) and (9). Equation (9) is a parabola in (y, x)-space with second derivative with respect to y equal to $(2(1+a)d)/((1+a-b)(-1+d))$. Equation (8) has second derivative equal to $-2(bd)^2/(1+b(1-d(1-y)))^3$. To complete the claim, we show that the second derivatives are never equal, so that the difference between the curves is strictly convex or concave, and thus has at most two roots. Equating the second derivatives and solving for a produces a solution $a = h(b, y, d)$ that is easily verified to be continuous and equal to -1 when, e.g., $d = 0$ and $d = 1$. Solving $h(b, y, d) = 0$ for y, it is easy to see that there is no solution for $y > 0$. Thus there are no values of $a > 0$ and $y > 0$ such that (8) and (9) have equal second derivatives.

We have now proved item 2 of Proposition 1.

As $d \to 1$, there are at least 2 equilibria. This can be verified from the limiting shape of Γ, given above, in particular the fact that the limiting upper boundary of Γ is $2/(b+q(1+a-b))$, along with the fact that the limiting shape of $p_t(q)$ satisfies $p_t(0) = 1/b < 2/b$.

B.4 Uniqueness of stable equilibrium

It is obvious from the above that for any (a, b, d), $(q, p) = (0, 0)$ is an SSE.[19] We want to show that, if there exists another stable equilibrium with $q > 0$, there is a unique such one.

Stationary equilibria occur when (i) the boundary of Γ intersects $t(q)$ or (ii) $(q, p) = (1, 1) \in \Gamma$.

[19] If there is an SSE with $q = 0$, then it must be that $p = 0$, and if there is an SSE with $p = 0$, it must be that $q = 0$.

Stability is captured more easily by solving $V = 0$ for q. This produces two solutions for q, call them $q_2(p) \leq q_1(p)$, that are an equivalent representation of the boundary of Γ. An SSE is a stationary equilibrium such that $q_1(p)$ intersects $t(q)$ or (ii) $(q, p) = (1, 1)$ lies in the interior of Γ.

Define q_0 as the solution to $p_t(q) = 1$. Define (q_β, p_β) by $q_\beta = \arg\min_q p_2(q)$ and $p_\beta = p_2(q_\beta)$, i.e., (q_β, p_β) is the lowest point of the boundary of Γ in (q, p)-space. Uniqueness of SSE is thus guaranteed by the following

LEMMA **2** $q_\beta > q_0$.

Proof.

Notice that Γ is null for sufficiently small δ and that, given (a, b) there is a smallest δ such that Γ is non-null. We show two facts. First, as Γ increases with d, it intersects the line $q = q_0$ from the right, so that $q_\beta > q_0$ at that d. Second, as Γ increases further, it does so in a way such that $q_\beta > q_0$ remains true for all $d < 1$.

Because p_2 is single-peaked, in order to prove these facts we show that $p_2(q)$ is decreasing in q at $q = q_0$ whenever d is large enough that $p_2(q_0)$ is defined. It is sufficient to show that the same properties hold for some $q^* > q_0$, again because of the fact that $p_2(q)$ is single-peaked. In what follows, using such an argument, we construct a set of exhaustive cases that prove $p_2(q)$ is decreasing at q_0 whenever it is defined.

From equation (3), note that $q_0 < b/(1 + b)$.

Define $s(a, b, d) = \frac{\partial p_2}{\partial q}|_{q=\frac{b}{1+b}}$. It is easy but very messy to write s explicitly.

We want to show that s is negative whenever it is defined.

It is easy to check that $s(a, b, 1) = 0$ and $\frac{\partial s}{\partial \delta}|_{\delta=1} = (1+b)^2/(2b) > 0$. Thus $s < 0$ for large δ.

By inspection, s has a vertical asymptote when

$$4a(1 + 5b)(-1 + d)d + 4d(d + b(-4 + 5d)) + a^2(-1 + d)(-1 + d + b(-1 + 5d)) = 0.$$

This equation is quadratic in d with both roots in the unit interval. It is the larger root that is of interest, since it is at that value of d above which Γ intersects the line $q = b/(1+b)$.

Denoting by $\hat{d}$ the larger root, we want to know that as $d \to \hat{d}$ from above, $s(a, b, d) \to -\infty$. Once that is proven, we know that, in fact, s is always negative for $\hat{d} < d < 1$. The

reason is the following: one can write s in the form $s = F * \sqrt{G} - H$, where F is linear in d and G and H are both quadratic in d. The roots of s must satisfy $F^2 G = H^2$, which is a degree-4 polynomial in d. Call this polynomial s^2. The leading coefficient of s^2 is $-(2+a)^6 b^2 < 0$. It is easy to verify that $s^2(a,b,0) = s^2(a,b,1) = 0$, and that s^2 is decreasing in d at $d = 0$. Thus, s^2 must have a root that is less than 0, and so it has at most one root strictly between 0 and 1. In particular, as d decreases from 1, it is impossible that s becomes first negative, then positive, and then asymptotes to $-\infty$.

We return now to determining the sign of $s(a,b,d)$ as d approaches the critical value from above. Simplifying, this reduces to the sign of

$$5\left(a^2-4\right)b^3 - (a+2)b\left(a^2(X-1) - 3a + 2X + 2\right) + (a+2)b^2\left(a^3 + 3a^2 + 12a - 5X - 4\right) - (a+2)X, \tag{10}$$

where

$$X = \sqrt{b\left(a^2 b + 6ab + 2a + 4b\right)}.$$

We want to show that expression (10) is negative. We use the fact that expression (10) is decreasing in X.

Case 1: $b < \frac{2}{5}a$. We have $X^2 = b\left(a^2 b + 6ab + 2a + 4b\right) > b^2(a+3)^2$, using the assumption that $b < \frac{2}{5}a$. Substituting, therefore, $b(a+3)$ for X in expression (10), and dividing by $(a+2)b$, we obtain

$$S_1 = a^2 + 2a(5b+1) - 5\left(5b^2 + 2b + 1\right). \tag{11}$$

Differentiating S_1 with respect to b produces $10(a-b-1) - 40b < 0$ provided $a - b < 1$, as we assume throughout. Thus S_1 is maximized when b takes its smallest value of $\max\{0, a-1\}$.

Case 1 (i): $a < 1$ and we evaluate S_1 at $b = 0$, obtaining $a^2 + 2a - 5$ which is negative for $a < 1$, as desired.

Case 1 (ii): $a > 1$ and we evaluate S_1 at $b = a - 1$, obtaining $-14a^2 + 32a - 20$ which is negative for $a > 1$, as desired.

Case 2: $b > \frac{2}{5}a$.

Given only that $a, b > 0$, we have that $X > (a+2)b$. Substituting, therefore, $(a+2)b$ for X into (10) yields an upper bound of

$$S_2 = a^2(b+1) + 2a(5b+1) - 4(5b^2 + 2b + 1). \tag{12}$$

Clearly S_2 is increasing in a, so it can be bounded above by taking maximal values of a.

Case 2 (i): $b < 0.62$.

We use the fact that $a < \frac{5}{2}b$ to substitue into S_2, yielding $\frac{25}{4}b^3 + \frac{45}{4}b^2 - 3b - 4$ which is negative in the claimed range.

Case 2 (ii): $1 < b < 5.82$.

We use the fact that $a < b + 1$ (assumed throughout) to substitue into S_2, yielding $b^3 - 7b^2 + 7b - 1$ which is easily verified to be negative in the claimed range.

Case 2 (iii): $5.82 < b$. We now use the fact that $b/(b+1) < 1$ and repeat the argument above, evaluating the slope of p_2 at $q = 1$. Define $\bar{s}(a, b, d) = \frac{\partial p_2}{\partial q}|_{q=1}$. We compute $2(a + 1)^2 d\bar{s}(a, b, d)$ as

$$(a+1)\left(\frac{a^2\left(-7d^2+8d-1\right)+a(d-1)(5bd-b-14d-2)+b\left(5d^2-4d-1\right)-7d^2+4d-1}{\sqrt{4(a+1)^2(d-1)d+(a(d-1)+d+1)^2}}+ad-a\right.$$
$$-(2a-b+2)\left(ad-\sqrt{4(a+1)^2(d-1)d+(a(d-1)+d\right.$$

It is easy to verify, parallel to the case with s, that $\bar{s}(a, b, 1) = 0$ and that $\frac{\partial \bar{s}}{\partial d} = b/2 > 0$.

The critical value of d for $\bar{s}$ is $d^* = \frac{1+4a+3a^2+2\sqrt{-1+2a+8a^2+6a^3+a^4}}{5(1+2a+a^2)}$. The sign of $\bar{s}(a, b, d^*)$ as $d \to d^*$ from above is determined by

$$4 + a^3 + 3\sqrt{-1+4a+a^2} - a^2\sqrt{-1+4a+a^2} - 5\left(3+\sqrt{-1+4a+a^2}\right)b$$
$$+a\left(3+2\sqrt{-1+4a+a^2}+5b\right). \tag{13}$$

Expression (13) is linear and decreasing in b, and so is negative for sufficiently large b, given a value of a. In particular, (13) is negative whenever

$$b > \frac{(1+a)\left(4+a^2+3\sqrt{-1+4a+a^2}-a\left(1+\sqrt{-1+4a+a^2}\right)\right)}{5\left(3-a+\sqrt{-1+4a+a^2}\right)} \equiv \bar{b}(a). \tag{14}$$

It is readily verified that $\bar{b}$ is increasing. To complete this part, we must show that (14) is satisfied for all $b > 5.82$ and $a < b + 1$ (the constraint that $a < \frac{5}{2}b$ is not binding for large b). Since $\bar{b}$ is increasing, a sufficient condition is that $b > \bar{b}(b + 1)$ for all $b > 5.82$. It is straightforward to verify that, in fact, this condition is met for all $b > \sqrt{21} - 3 \approx 1.58$.

Case 2 (iv): $.62 < b < 1$. We proceed in two steps. Step 1 completes the case for an upper bound on a. Step 2 shows that the bound is increasing in a, so that Step 1 is in fact the worst case scenario.

Step 1: We show the desired result, that the derivate of $p_2(q)$ diverges to $-\infty$ at the critical δ, under an upper bound on a. Since $b < 1$, $q_0 < b/(1+b) < 1/2$. Thus it is enough to examine $\bar{s}'(a,b,d) = \frac{\partial p_2}{\partial q}|_{q=1/2}$ as d approaches the critical value. Simplifying yields $9(1+a+b)\bar{s}'(a,b,d)$ equal to

$$3+a^3-33b+13b^2+b^3+a^2(1+3b)+a\left(3+14b+3b^2\right)-\left(-3+a^2+2a(-1+b)+10b+b^2\right)X', \tag{15}$$

where

$$X' = \sqrt{-3+a^2+6b+b^2+2a(3+b)}.$$

The above expression is decreasing in X' in the relevant parameter range and so is bounded above by taking a lower bound for X'. X' is clearly increasing in b and so is bounded below by setting $b = \frac{2}{5}a$. The resulting expression is $(X')^2 = -3 + \frac{42a}{5} + \frac{49a^2}{25}$. We are here concerned with a maximal value of a, which, in this case, falls in $[5/2 * .62, 2 * 1] = [1.55, 2]$. It is easily verified that, in this range, a lower bound for that expression of $(X')^2$ is $(\frac{7}{5}a + \frac{5}{3})^2$. Substituting this last expression for X' gives us an upper bound for the expression of interest:

$$-6a^3 + a^2(32+3b) + 2a\left(79-25b+12b^2\right) + 5\left(24-149b+34b^2+3b^3\right). \tag{16}$$

Now, when $b < 2/3$ we have $a < \frac{5}{2}b$ and when $b > 2/3$ we have $a < b+1$. Making those two substitutions into the above yields $5\left(24-70b+49b^2\right)$ and $4\left(76-147b+41b^2+9b^3\right)$, respectively, each of which is easily verified to be negative for the relevant range of b.

Step 2: We show that the maximal value of a is the worst case scenario by proving that (15) is increasing in a.

Taking the derivative with respect to a and simplifying, we get:

$$-12b(3+a+b)+4(1+a+b)(3+a+b)-(1+a+b)^2(3+a+b)+$$
$$4(-3+a^2+2a(3+b)+b(6+b))-2(1+a+b)(-3+a^2+2a(3+b)+b(6+b))+$$
$$(3+3a^2+b(14+3b)+a(2+6b))\sqrt{-3+a^2+2a(3+b)+b(6+b)}.$$

Call this $y(a,b)$. Note that the coefficient of the radical is positive. To show that $y(a,b)$ is always positive, we will substitute a lower bound for this radical and show the resulting polynomial is always positive.

Write $Y=-3+a^2+2a(3+b)+b(6+b)$. We claim that $Y>(a+1.69b)^2$ over the range $b\in[0.62,1]$ and $a\in[0,2]$. This inequality is equivalent to

$$-(1.69^2-1)b^2+6a-(1.38)ab+6b-3>0.$$

Noting that the LHS is increasing in a (as $1.38b<6$); it is minimized at $a=0$. We are left with a quadratic equation in b, and one can easily verify that the range $[0.62,1]$ lies between the two roots of this quadratic, and hence the LHS is positive over that range (since the coefficient of b^2 is negative).

We can therefore substitute $(a+1.69b)$ for $\sqrt{Y}$ in $y(a,b)$. Simplifying then gives the polynomial

$$3.-3.93b+0.66b^2+2.07b^3+a^2(-9.+2.07b)+a(30.-16.62b+4.14b^2).$$

One can then easily verify that this polynomial is positive over the range $b\in[0.62,1]$ and $a\in[0,2]$. We conclude that $y(a,b)$ is positive in this range, as we required. □

B.5 Final pieces

Theorem 1 is now proved.

To complete the proof of Proposition 1, we remark that the second and fourth items are clear from the shapes of $t(q)$ and Γ already discussed. Next, notice that if there are four SSE, then the largest two involve $p=1$, since at most two can be interior. In this case it must be that exactly one of them involves q_1 and is stable.

Using the characterization of stable SSE in the proof above, Proposition 2 follows easily. It uses only the limiting shape of Γ already derived and the fact that as $d\to 1$, $q_0\to\max\{\frac{b-1}{b},0\}$.

References

[1] Jeremy Blackburn, Ramanuja Simha, Nicolas Kourtellis, and Xiang Zuo. Cheaters in the Steam Community Gaming Social Network. *Arxiv preprint arXiv:*, 2011.

[2] Simon Board. Relational contracts and the value of loyalty. *American Economic Review*, 2010.

[3] Simon Board and Moritz Meyer-ter vehn. Relational Contracts in Competitive Labor Markets. pages 1–39, 2011.

[4] L Dall'Asta, Matteo Marsili, and P Pin. Collaboration in Social Networks. *Arxiv preprint arXiv:1104.2026*, (5):1–18, 2011.

[5] S. Datta. Building trust. Mimeo, Indian Statistical Institute, 1993.

[6] G. Ellison. Cooperation in the prisoner's dilemma with anonymous random matching. *The Review of Economic Studies*, 61(3):567–588, July 1994.

[7] T. FUJIWARA-GREVE and M. OKUNO-FUJIWARA. Voluntarily Separable Repeated Prisoner's Dilemma. *Review of Economic Studies*, 76(3):993–1021, 2009.

[8] P. Ghosh and D. Ray. Cooperation in community interaction without information flows. *The Review of Economic Studies*, 63(3):491–519, July 1996.

[9] M. Jackson, T. Rodr\'\iguez Barraquer, and X. Tan. Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*, 2011.

[10] M. Jackson and B.W. Rogers. Meeting Strangers and Friends of Friends: How Random are Social Networks? *The American Economic Review*, 97(3), 2007.

[11] Matthew O. Jackson and Alison Watts. Social games: Matching and the play of finitely repeated games. *Games and Economic Behavior*, 70(1):170–191, September 2010.

[12] M. Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59:63–80, 1992.

[13] D Krackhardt. Social Networks and the Liability of Newness for Managers. In C. Cooper and M. Rousseau, editors, *Trends in Organizational Behavior, Volume 3*, chapter 9, pages 159–173. John Wiley & Sons Ltd, New York, NY, 1996.

[14] R. Kranton. The formation of cooperative relationships. *Journal of Law, Economics, and Organization*, 12(1):214–233, April 1996.

[15] J Levin. Multilateral contracting and the employment relationship. *The Quarterly Journal of Economics*, (August):1075–1104, 2002.

[16] W.B. MacLeod and J.M. Malcomson. Motivation and markets. *American Economic Review*, pages 388–411, 1998.

[17] G. Ramey and J. Watson. Bilateral trade and opportunism in a matching market. *Contributions to Theoretical Economics*, 1(1), 2001. Article 3.

[18] A Riedl and A Ule. Exclusion and cooperation in social network experiments. *Unpublished Paper, CREED, University*, 2002.

[19] C. Shapiro and J.E. Stiglitz. Equilibrium unemployment as a worker discipline device. *American Economic Review*, 74(3):433–444, June 1984.

[20] J. Sobel. A theory of credibility. *Review of Economic Studies*, 52:557–573, 1985.

[21] J. Sobel. Can we trust social capital? *Journal of Economic Literature*, XL:139–154, March 2002.

[22] F. Vega-Redondo. Building up social capital in a changing world. *Journal of Economic Dynamics and Control*, 30(11):2305–2338, November 2006.

[23] J. Watson. Starting small and renegotiation. *Journal of Economic Theory*, 85:52–90, 1999.

[24] J. Watson. Starting small and commitment. *Games and Economic Behavior*, 38:176–199, 2002.