UC Irvine

UC Irvine Previously Published Works

Title

A dynamic model of social network formation.

Permalink

<https://escholarship.org/uc/item/9652h1dp>

Journal

Proceedings of the National Academy of Sciences of the United States of America, 97(16)

ISSN

0027-8424

Authors

Skyrms, B
Pemantle, R

Publication Date

2000-08-01

DOI

10.1073/pnas.97.16.9340

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

A dynamic model of social network formation

Brian Skyrms^{*†} and Robin Pemantle[‡]

^{*}School of Social Sciences, University of California, Irvine, CA 92607; and [‡]Department of Mathematics, Ohio State University, Columbus, OH 43210

This contribution is part of the special series of Inaugural Articles by members of the National Academy of Sciences elected on April 27, 1999.

Contributed by Brian Skyrms, May 9, 2000

We consider a dynamic social network model in which agents play repeated games in pairings determined by a stochastically evolving social network. Individual agents begin to interact at random, with the interactions modeled as games. The game payoffs determine which interactions are reinforced, and the network structure emerges as a consequence of the dynamics of the agents' learning behavior. We study this in a variety of game-theoretic conditions and show that the behavior is complex and sometimes dissimilar to behavior in the absence of structural dynamics. We argue that modeling network structure as dynamic increases realism without rendering the problem of analysis intractable.

Pairs from among a population of 10 individuals interact repeatedly. Perhaps they are cooperating to hunt stags and rabbits, or coordinating on which concert to attend together; perhaps they are involved in the somewhat more antagonistic situation of bargaining to split a fixed payoff, or attempting to escape the undesirable but compelling equilibrium of a Prisoner's Dilemma. As time progresses, the players adapt their strategies, perhaps incorporating randomness in their decision rules, to suit their environment. But they may also exert control over their environment. The players may have choice over the pairings but not perfect information about the other players. They may improve their lot in two different ways. A child who is being bullied learns either to fight better or to run away. Similarly, a player who obtains unsatisfactory results may choose either to change strategies or to change associates. Regardless of whether the interactions are mostly cooperative or mostly antagonistic, it is natural and desirable to allow evolution of the social network (the propensity for each pair to interact) as well as the individuals' strategies.

We build a model that incorporates both of these modes of evolution. The idea is simple.

(*)

Individual agents begin to interact at random. The interactions are modeled as games. The game payoffs determine which interactions are reinforced, and the social network structure emerges as a consequence of the dynamics of the agents' learning behavior.

As the details of the specific game and the reinforcement dynamics vary, we then obtain a class of models. In this paper, we treat some simple reinforcement dynamics, which may serve as a base for future investigation.

The idea of simultaneous evolution of strategy and social network appears to be almost completely unexplored. Indeed, the most thoroughly studied models of evolutionary game theory assume *mean-field* interactions, where each individual is always equally likely to interact with each other. Standard treatments of evolutionary game dynamics (1, 2) operate entirely in this paradigm. This is due, to a large extent, to considerations of theoretical tractability of the model. Models have been introduced that allow the agents some control over their choice of partner (3), but the control is still exerted in a mean-field setting: one chooses between the present partner and a new pick at random from the whole population.

Evolutionary biologists know that evolutionary dynamics can be affected by nonrandom encounters or population structure,

as in Sewall Wright's models of assortative mating (4). Wright (5) already realized that positive correlation of encounters could provide an account of evolution of altruism. Thus, the need for social network models has been long recognized.

When the social network is modeled, it is almost always static.[§] Interactions, for example, may be posited to occur only between players whose locations are close, according to some given spatial data. Biological models in which encounters are governed by spatial structure have become increasingly frequent in the 1990s; see, for example, the work of Durrett, Levin, and Neuhauser (7–9). A similar hypothesis of spatial structure, in a game theory context, arises in ref. 10. Here, technology from statistical mechanics is adapted to the analysis of games whose interactions take place between neighbors in a grid.

A number of recent investigations by game theorists, some directly inspired by biological models, have shown that the dynamics of strategic interaction can be strikingly different if interaction is governed by some spatial structure, or more generally, some graph structure (11–13). For instance, one-shot Prisoner's Dilemma games played with neighbors on a circle or torus allows cooperation to evolve in a way that the random encounter model does not. The spatial or graph structure can be important to determine which equilibria are possible, whether repeated interactions can be expected to converge to equilibrium, and, if so, how quickly convergence takes place (14).

Because the outcome of a repeated game may vary with the choice of network model, it is important to get the network model right. Further progress in the theory of games and adaptive strategies would be greatly enhanced by a theory of networks of social interaction. In particular, it would be desirable to have a framework within which models may be developed that are both tractable and plausible as a mechanism governing interactions among a population of agents seeking to improve their lot.

When the network changes much more slowly than do the strategies of individuals, it is reasonable to model the social network by a structure that is fixed, though possibly random. The question of realistically modeling the randomness in such a case is taken up in a number of papers, of which a recent and well known example is the "small world" model (15). In the other extreme (16–18), evolution of social structure is modeled by agents moving on a fixed graph in the absence of strategy dynamics.

In the general case, however, interaction structures are fluid and evolve in tandem with strategy. What is required here is a dynamics of interaction structure to model how social networks are formed and modified. We distinguish this *structure dynamics* from the *strategic dynamics* by which individuals change their individual behaviors or strategies.

In this paper, we introduce a simple, additive model for structure dynamics, and we explore the resulting system under several conditions: with or without discounting of the past, with or without added noise, and in the presence or absence of strategic dynamics. Common to all our models is a stochastic evolution

[†]To whom reprint requests should be addressed. E-mail: bskyrms@uci.edu.

[§]An exception, perhaps, is a preprint we have recently encountered by Jackson and Watts (6).

from a (usually symmetric) initial state. Individuals in a population start out choosing whom to interact with at random and then modify their choices according to how their choice is reinforced, and then the process is repeated. An infinite variety of such models is possible. We will consider only a few basic models, meant to illustrate that rigorous results on structure dynamics are not out of reach, and that further inquiry will be profitable.

We first consider a baseline case of uniform reinforcement. Here, any choice of partner is reinforced as strongly as any alternative choice would have been. In other words, the interaction game between any pair of players always produces a constant reward or punishment. One might expect that such cases would not lead to interesting dynamics, but that is far from the truth. We show both by computer simulation and analytically how structure emerges spontaneously even in these cases. Because the strategic dynamics here are trivial, the baseline case is intended mostly as a building block on which more interesting strategic dynamics are to be grafted. We note, however, that the constant reward game is not completely unreasonable. Studies have shown that in the absence of other environmental attributes, sheer familiarity brings about positive attitudinal change (19). In fact, an abstract model of network evolution under uniform positive reweighting has appeared before under the name of “Reinforced Random Walk” (D. Coppersmith and P. Diaconis, unpublished work).

Next, we move to the case where players of different types play a nontrivial game and are reinforced by the payoffs of the game. Here, we examine the coevolution of behavior and structure when the structural dynamics and strategic dynamics are both operative. The relative speeds of structural dynamics and strategic dynamics affect which equilibrium is selected in the game. In particular, this can determine whether the risk-dominant or payoff-dominant equilibrium is selected.

Making Friends: A Baseline Model of Uniform Reinforcement

Friends I: Asymmetric Weights. Each morning, each agent goes out to visit some other agent. The choice of whom to visit is made by chance, with the chances being determined by the relative *weights* each agent has assigned to the others. For this purpose, agent number i has a vector of weights $\langle w_{i1}, \dots, w_{in} \rangle$ that she assigns to other players (assume $w_{ii} = 0$). Then she visits agent j with probability

$$\text{Prob}(\text{agent } i \text{ visits } j) = \frac{w_{ij}}{\sum_k w_{ik}}. \quad [1]$$

Here we are interested in a symmetric baseline model, so we will assume that all initial weights are 1. Initially, for all agents, all possible visits are equiprobable.

Every agent is treated nicely on her visit and all are treated equally nicely. They each get a reinforcement of 1. Each agent then updates her weight vector by adding 1 to the weight associated with the agent that she visited. Her probabilities for the next round of visits are modified accordingly. At each stage, we have a matrix p_{ij} of probabilities for i to visit j . Do these probabilities converge, and if so to what?

Given all the symmetry built into the starting point and the reinforcement, it is perhaps surprising that all sorts of structures emerge. Here is a description of a simulated sample run of length 1,000. The probabilities, to two decimal places, seem to converge after a few hundred rounds of visits, to a matrix that is anything but uniform (and to a different matrix each time the process is run from the initial, symmetric weights). There is one agent, A, who visits another agent, B, more than half the time. There is no reciprocation, so this has no bearing on how often B visits A, and in fact most agents will not visit any one agent more than a third of the time.

In the analysis section, we show that this outcome is typical.

THEOREM 1. *The probability matrix for Friends I with n players will converge to a random limit p as time goes to infinity. The distribution of the limit is that the rows of p are independent, each having Dirichlet distribution (ignoring the zero entry on the diagonal) whose parameters are $n - 1$ ones.*

Thus we see spontaneous emergence of structure. This type of simple model has been used before in the economics literature to explain the stabilization of market shares at seemingly random equilibria, due to random reinforcement in the early phases of growth of an industry (20). We remark that the choices made by each agent are independent of the choices made by each other agent, so the social aspect of the model is somewhat degenerate and the model may be viewed as a model of individual choice. Nevertheless, it fits our definition of social network model in that it gives a probabilistic structure to interactions; one may then extend the model so the interactions are nontrivial games.

Friends II: Symmetrized Reinforcement. Suppose now that the interaction is as pleasant to the host as the visitor. Thus when agent i visits agent j , we add 1 to both w_{ij} and w_{ji} . A typical outcome for 10 agents after 1,000 rounds of visits looks similar to the table for Friends I, except that the entries are nearly symmetric. There are, however, subtle differences that may cause the two models to act very differently when strategic dynamics are introduced. To see these differences, we describe what is typically observed after 10 runs of a simulation of Friends II to time 1,000 for a set of three agents, this being the minimum population size for which structural dynamics are interesting. What we see typically is one or two runs in which each player visits are split evenly (to two decimal places) between the others. We see another several runs that are close to this. We see one run or so in which two agents nearly always visit the third agent, which splits its time among the other two. The remaining runs give something between these extreme outcomes.

What may not be apparent from such data is that the limiting weights for Friends II are always $1/2$. Only a small fraction of sample outcomes decisively exhibit the proven limiting behavior. The data, in other words, show that after 1,000 iterations, the weights may still be far from their limiting values; when this is the case, one of the three agents is largely ignored by the other two and visits each of the other two herself equally often. Because the lifetime of many adaptive games is 1,000 rounds or fewer, we see that limiting behavior may not be a good guide to behavior of the system on time scales we are interested in. The analysis section discusses both limiting results for this model and finite time behavior. When the population size is more than 3, the weights will always converge, but the limit is random and restricted to the subspace of symmetric matrices. Again, convergence of the weights to their limiting values is slower than in the nonreciprocal game of Friends I.

THEOREM 2. *The probability matrix p_{ij} for Friends II with n players converges to random limit p as time goes to infinity. If $n = 3$, the limit is the matrix all of whose off-diagonal entries are $1/2$. In general, the limit may be any symmetric matrix whose rows sum to 1; that is, the closed support of the random limit is the entire subspace of symmetric stochastic matrices.*

Analysis of Friends I and II. To fit this in the framework of (*), construct the following degenerate games. Each of the two players has only one strategy, and the payoff matrix is as follows.

Friends I	Host	Friends II	Host
Visitor	(1, 0)	Visitor	(1, 1)

The weights w_{ij} are initialized to 1 for $i \neq j$, and are then updated according to

$$w_{ij}(t+1) = w_{ij}(t) + u(i, j; t), \quad [2]$$

where $w_{ij}(t)$ is the weight agent i gives to agent j at time t and $u(i, j; t)$ is the payoff of the game played at time t between visitor i and host j (and zero if this visit did not occur at time t). This, together with specification of the visitation probabilities in Eq. 1, defines the model. Changing the initial weights does not affect the qualitative behavior of any model, so there is no need to vary the initialization.

For Friends I, the updating of the weights for any one agent is the same as a Pólya urn process (21). Each agent can be thought of as having an urn with balls of $n - 1$ colors, one color representing each other agent. Initially there is one ball of each color in the urn. The agent picks a ball at random, indicating whom she should visit, then returns it to the urn along with an extra ball of the same color. The urns belonging to different agents are statistically independent.

The analysis of this process is well known (ref. 22, chapter 4). It is easy to show that the sequence of draws for each agent is *exchangeable*, that is, permuting a sequence does not change its probability. Hence by the de Finetti representation theorem, the random sequence of draws from an urn is equivalent to a mixture of multinomial processes, that is, of sequences of independent draws. The mixing measure is easily seen to be Dirichlet. Consequently, the visiting probabilities converge with probability one, but they can converge to anything. That they converge to the uniform vector, where each agent has equal probability to visit each other, has prior probability zero.

Furthermore, convergence to the limiting probability matrix is quite rapid. Let $p(t)$ denote the matrix whose (i, j) -entry is $p_{ij}(t)$. Then exchangeability implies that, conditional on the limit matrix $p = \lim_{t \rightarrow \infty} p(t)$, the sequence of visits is a sequence of independent, identically distributed draws from the limit distribution. Thus at time t , the central limit theorem implies that $p(t) - p$ is $t^{-1/2}$ times a multivariate normal.

For Friends II, exchangeability fails. This is not surprising, because the property of exchangeability is not very robust. More surprising, however, is that the sequence of probability matrices $p(t)$ does not form a martingale. To explain this terminology, let $\mathbf{E}_t$ denote the expectation conditioned on the values at time t . A simple computation shows that for Friends I, the expected value of $p_{ij}(t + 1)$ conditioned on the time t value is equal to $p_{ij}(t)$: because w_{ij} increases only when i visits j , we have

$$\begin{aligned} \mathbf{E}_t p_{ij}(t + 1) &= \mathbf{E}_t \sum_{k=1}^n p_{ik}(t) \frac{w_{ij} + \delta_{jk}}{1 + \sum_{l=1}^n w_{il}(t)} \\ &= \frac{w_{ij}(t) + p_{ij}(t)}{1 + \sum_{l=1}^n w_{il}(t)} \\ &= p_{ij}(t). \end{aligned}$$

Even without exchangeability, the martingale convergence theorem (ref. 23, section 4.2) implies convergence of the quantities p_{ij} , though it says very little about the limit.

For Friends II, a complete analysis may be obtained (R.P. and B.S., unpublished work). Here is an outline of what is found there. A computation similar to the one for Friends I shows that

$$\mathbf{E}_t p(t + 1) = p(t) + \frac{1}{t} F(p(t)),$$

where F is a certain function on symmetric n by n matrices. In other words, the random sequence of matrices $\{p(t) : t = 1, 2, \dots\}$ is a *stochastic approximation* in the sense of Robbins and Monro (24), driven by the vector field F . General results of refs. 25 and 26 now imply that $p(t)$ converges to the set where F vanishes. To show that $p(t)$ always converges to a single point, Pemantle and Skyrms (unpublished work) compute a Lyapunov function for F , that is, a function V for which $\nabla V \cdot F < 0$ with equality only when $F = 0$. This, together with an efficiency inequality (bounding the angle between f and ∇V away from

ninety degrees), establish convergence of p . The remainder of *Theorem 2* is then established by showing the only stable zeros of the vector field F are the symmetric matrices with row sums all equal to 1, and that the possible limit points of $p(t)$ are exactly the stable equilibria of the flow determined by F .

Determination of the rate of convergence of $p(t)$ to its limit is somewhat different in this case. Because of the presence of unstable equilibria from the flow determined by F , there is a possibility of being stuck near one of these equilibria for a long time before eventually following the flow to one of the stable equilibria. For the three player game, the unstable equilibria are the following three matrices:

$$\begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & 0 & 1 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ 1 & 0 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix}.$$

These correspond to cases where one of the three agents is entirely ignored and splits her visits equally between the other two. The probability that $p(t)$ is within ϵ of one of these traps is roughly $3\epsilon t^{-1/3}$, so with $t = 1,000$, we find a reasonably high probability that $p(1,000)$ is not near the uniform probability matrix but is instead still near one of the unstable equilibria. This persists with reasonable probability well beyond $t = 10^6$. For greater population sizes, similar phenomena apply. Convergence to the invariant set is relatively slow. However, for large populations, say 20 or more, another phenomenon takes place. The portion of the space of possible p matrices that are within ϵ of the possible limits goes to 1; this is known as the concentration of measure phenomenon (27). Thus it becomes very unlikely to get stuck initially far away from the limit, simply because the initial randomness will very likely lead to a point very near a possible limit. Thus for large populations, the dynamics appear very similar to the dynamics for Friends I.

Making Enemies

Let us change the “Making Friends” model in just one way. Instead of being rewarded, agents are punished; instead of uniformly positive interactions, we have uniformly negative ones:

Enemies I	Host	Enemies II	Host
Visitor	$(-1, 0)$	Visitor	$(-1, -1)$

Instead of interactions being reinforcing, we take them as inhibiting. The dynamics of inhibition might be modeled in a number of ways. Continuing to use the update Eq. 2 will not work because the weights will end up becoming negative and the visitation probabilities in Eq. 1 will be meaningless. In this section, we explore two other possible rules for updating the weights so as to inhibit past behavior. With negative reinforcement, it is easy to predict what will happen: the social network always becomes uniform, and the dynamics are not sensitive to the particular updating mechanism. Indeed, this is what happens. Because there are no surprises, and because this model is just a building block for a model with both structural and strategic dynamics, we keep the discussion brief.

The Transfer Model. Consider a three-player model with the following update rule on the weights. Initial weights are all positive integers. When i visits j , the weight w_{ij} is diminished by 1 and the weight w_{ik} , $k \neq i, j$, is increased by 1. This is equivalent to the Ehrenfest model of heat exchange between two bodies (28). In the original Ehrenfest model, there are two urns. A ball is drawn at random from among all balls in both urns and transferred to the other urn. The distribution of balls tends to the binomial distribution, where each ball is independently equally likely to be in either urn. In Making Enemies, with transfer

dynamics and three players, each player may be thought of as having such a pair of urns. The urns are independent.

Because the number of balls is fixed, an Ehrenfest urn is a Markov chain with a finite number of states, where the states consist of distributions over the two urns. For example, if there are only two balls, then there are three states, S_1 , S_2 and S_3 , corresponding to urn cardinalities of $(2, 0)$, $(1, 1)$, and $(0, 2)$. The transition matrix for this Markov chain is

$$\begin{pmatrix} 0 & 1 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 1 & 0 \end{pmatrix},$$

and the unique stationary vector is $(1/4, 1/2, 1/4)$. In contrast to the Pólya urn, we do not have convergence of the conditional probabilities of visits at each stage given the present: at any time, given the present composition, the probability of a given visit may be 0, $1/2$, or 1, depending on the composition of the urns belonging to the visitor. However, if the number of balls, N is large, approximately equal visiting probabilities are very likely in the following sense. The invariant distribution is binomial, which is concentrated around nearly even distributions when the number of balls is large. Thus, with high probability, no matter what the initial state, after roughly $N \log N/2$ steps (29), the composition of an urn with N balls will be close to a draw from a binomial distribution. The conditional probability of either of the two possible visits will therefore be close to $1/2$ and will tend to remain there with high probability. Kac (30) uses these properties to resolve the apparent paradoxes that beset Boltzmann's discussion of irreversibility in statistical mechanics.

The Resistance Model. The transfer model allows for a finite cumulative amount of negative reinforcement, and indeed yields a finite Markov chain. Let us explore a rather different model, termed the *resistance model*, in which negative payoffs generate resistance. Initially every choice has resistance 1. The magnitude of a negative payoff is added to its associated resistance, so the Eq. 2 becomes

$$w_{ij}(t+1) = w_{ij}(t) + |u(i, j; t)|.$$

In the case at hand, when all payoffs are negative, the probability of i visiting j is proportional to the reciprocal of the resistance:

$$p_{ij} = \text{Prob}(\text{agent } i \text{ visits } j) = \frac{1/w_{ij}}{\sum_{k=1}^n 1/w_{ik}}$$

with $1/w_{ii} = 0$ by convention. The dynamics of Enemies I and Enemies II under resistance dynamics are easy to describe.

THEOREM 3. *For Enemies I or Enemies II, from any initial conditions, the probability matrix $p(t)$ converges to the uniform probability matrix $\bar{p}$ where $\bar{p}_{ij} = 1/(n-1)$ for any $i \neq j$. The convergence is rapid: of order $N \log N$ if the initial resistances are of order N . The deviations from uniform obey a central limit theorem:*

$$t^{1/2}(p - \bar{p}) \rightarrow X$$

where X is a multivariate normal with covariance matrix of rank $n(n-1)$ in Enemies I and $n(n-1)/2$ in Enemies II. In other words, deviations from uniformity are independent normals, subject to the constraints of adding up to zero for each individual and, in the case of Enemies II, the constraints of symmetry.

The central limit theorem may be derived from a stronger, functional central limit theorem, linearizing the system near the uniform probability to see that the paths

$$t \mapsto N^{-1/2}(p(Nt) - \bar{p})$$

converge in distribution as $N \rightarrow \infty$ to a multivariate Ornstein-Uhlenbeck process. The rate of convergence follows from standard coupling arguments.

While uniform positive reinforcement breeds structure from unstructured initial conditions, uniform negative reinforcement evidently breeds uniformity even from structured initial conditions. It would appear, therefore, that the customary random encounter (mean-field) model is more suitable for Making Enemies than Making Friends.

A Better Model? We would like a model that allows for both positive and negative reinforcement. A natural choice is to let w_{ij} keep track of the log-likelihood for i to visit j , so that probability of i visiting j is given by

$$p_{ij} = \text{Prob}(\text{agent } i \text{ visits } j) = \frac{\exp(w_{ij})}{\sum_{k=1}^n \exp(w_{ik})}. \quad [3]$$

In the next section, we will see a property this rule has in common with rules that discount the past, namely that it leads to being trapped in a deterministic state where i always visits the same j .

QUESTION 1. *Is there a model incorporating both positive and negative reinforcement, that is realistic, tractable, and nontrapping?*

Perturbations of the Models

In this section, we add two features, noise and discounting, commonly used to create more realistic models. We examine the effects on social structure. In particular, these lead to varying degrees of subgroup formation.

Discounting the Past. In the foregoing models, a positive (or negative) payoff in the distant past contributes equally to the weight (or resistance) assigned to an edge as does a like payoff in the immediate past. This is implausible, both psychologically and methodologically. As a matter of psychology, memories fade. From the standpoint of inductive logic, it is not at all certain that the learner is dealing with stationary probabilities—indeed, in cases of prime interest, she is not. For this reason, recent experience may have a better chance of being a relevant guide to future action than the remote past.

A simple and standard way to modify the models to reflect this concern is to introduce discounting of the past. We will concentrate here on the models of Making Friends. After each interaction, we will now multiply the weights of the previous stage by a discount factor, d , between 0 and 1. The we add the undiscounted payoffs from the present interaction to get new weights. The modification of the dynamics has a dramatic effect on the Making Friends models.

For Friends I, it is immediately evident from simulations with $d = 0.9$, say, and 10 players, that the probabilities p_{ij} converge to 0 or 1. In other words, each individual ends up always visiting the same other individual.

In Friends II, simulations show the group breaking into pairs, with each member of a pair always visiting his or her “partner.” Which pairs form depends on the randomness in the early rounds of visits, but pairs always form. In fact, there are other possible limit states, but their frequency is low except at more extreme discount rates. The set of possible limit states may be described as follows. Some agents are grouped in pairs, each member of a pair always visiting the other. Other agents are grouped in *stars*. These are clusters of size at least three, in which one agent, called the *center*, visits each of the others with positive frequency, while the others always visit the center.

Analysis of Discounting the Past. It is worth giving a rigorous derivation of the above behavior, because it will shed some light on a defect in the most obvious log-likelihood model to incorporate positive and negative reinforcement. Our derivation highlights this, although the results for discounted Friends I may also be derived from a theorem of H. Rubin (see ref. 31, page 227).

THEOREM 4. *In Friends II with discount rate $d < 1$, there is always a partition into pairs and stars and a random time after which each member of a pair visits only the other member of the pair and each noncentral member of a star visits only the center. In Friends I, there is a random function f and a random time after which each player i always visits $f(i)$.*

Sketch of Proof: The analysis for Friends I is similar but easier, so we prove the statement only for Friends II. With each probability matrix p we associate a graph $G(p)$ as follows. The edge (i, j) is in the graph G if the probability $p_{ij} > \epsilon$, where $\epsilon < 1/(2n)$ is some fixed positive number. Among those graphs having at least one edge incident to each vertex, let S denote the minimal such graphs, that is, ones for which deleting any edge results in an isolated vertex. It is easy to see that S is the set $G(p)$ for all p satisfying the conclusion of the theorem.

The principle behind the analysis of discounted Friends is that the future behavior of p is largely determined by the present $G(p)$. In particular, we find a $\delta > 0$ such that from any state p , for each subgraph H of $G(p)$ such that $H \in S$, there is a probability at least δ^2 that for all sufficiently large t , $G(p(t)) = H$. We show this in two steps: (step *i*) with probability at least δ , there is some t for which $G(p(t)) = H$; (step *ii*) from any state p such that $G(p) = H$, there is probability at least δ that $G(p(t))$ is equal to H for all later times, t .

To see why step *i* is true, for $H \in S$, let f_H be any function on vertices of H for which each value $f(i)$ is a neighbor of i . Observe that there is a number k such that from any state p with $H \subseteq G(p)$, if each vertex i visits $f(i)$ for the next k rounds, then $G(p(k)) = H$. For each round of visits, this probability is at least ϵ^n , where n is the number of vertices, so taking $\delta \leq \epsilon^{kn}$ establishes (step *i*). For (step *ii*), it suffices to show that with probability δ each agent visits a neighbor in H at all later times. For each agent i , the sum over j not neighboring i in H of p_{ij} is at most $n\epsilon < 1/2$ by the definition of $G(p) = H$. After k rounds of visits where agents only visit their neighbors in H , this must decrease to at most $(1/2)d^k$. Thus the probability of N rounds of visits only to neighbors in H is at least

$$\prod_{k=0}^{N-1} \left(1 - \frac{1}{2}d^k\right)^n.$$

Sending N to infinity yields a convergent infinite product, since $(1/2)d^k$ is summable. Taking δ to be less than the infinite product proves (step *ii*).

With steps *i* and *ii*, the rest is a standard tail argument. The constraints on evolution are such that $G(p(t))$ always contains at least one graph in S . As long as it contains more than one graph in S , there is always a probability of at least δ of permanently settling into each one. Thus, with probability 1, eventually $G(p(t))$ is equal to some $H \in S$ for all future times. This is equivalent to the conclusion of the theorem. **QED**

Remark: It is actually shown that in (step *ii*), if we choose ϵ sufficiently small, we can choose δ arbitrarily close to 1.

We now also see why the log-likelihood rule (3) leads to fixation of a degenerate structure. Under these dynamics, an equivalent phenomenon occurs to step *i* in the proof of *Theorem 4*. For a pair (i, j) whose interaction has a positive mean, if the pair plays repeatedly, we will see $w_{ij}(t)/t \rightarrow \mu > 0$. The probability the i will ever switch partners, once having tried j a few times is at most on the order of $\sum_{k=0}^{\infty} B \exp(-k\mu)$, where $B = \exp(\sum_{l \neq j} w_{il})$. From here it is easy to construct an argument parallel to the proof of *Theorem 4*, to show that in presence of a game with positive mean payoff, discounted structural dynamics lead with probability 1 to fixation at a pairing.

Introduction of Noise. A common feature in models of adaptation is the introduction of noise: a small chance of a behavior other than the one chosen by the dynamical equation for the model.

This may stem from an agent's uncertainty, from agent error, or from circumstances beyond an agent's control. Alternatively, an agent may purposefully add noise to her strategy in order to avoid becoming wedded to a less than optimally efficient strategy or structure.

From a methodological point of view, noise that does not go to zero with time transforms the model into an ergodic Markov chain. No state is then trapping. To the extent that the trapping states produced by discounting or linear log-likelihood are unrealistic, we may hope to mitigate the problem by adding a noise component. Because dynamics with a noise term do not lead to a single state, the outcome is usually phrased in terms of *stochastically stable states* (32). A state is termed stochastically stable if the chance of finding the system near that state does not go to zero as the magnitude of the noise term goes to zero.

Neither discounting nor noise will affect the limiting behavior of Making Enemies. For Making Friends, let us modify the probability rule (1) so that in the n -player game, the probability of i visiting j is now some fixed positive number $\epsilon/(n-1)$, plus $(1-\epsilon)$ times what it was before:

$$p_{ij} = \frac{\epsilon}{n-1} + (1-\epsilon) \frac{w_{ij}}{\sum_k w_{ik}}.$$

The effect of this is to push the system by ϵ toward the uniform point $\bar{p}$. Neither Friends I nor Friends II is now a martingale, and the stable set of each is reduced to the single point $\bar{p}$. Because this is true at any noise level $\epsilon > 0$, we see that there is only one asymptotically stable point. Because the qualitative outcome is sensitive to the existence of a noise term, it is incumbent to ask with regard to specific models whether a noise term is natural and realistic.

Noise and Discounting. In the presence of a discount $d < 1$ and a noise term $\epsilon > 0$, if $1-d$ is much smaller than ϵ , then the discount is so low that the noise term wipes out any effect the discounting might have had. In the other case, where d is held fixed and ϵ tends to zero, we may ask about the asymptotically stable states of system with past discounting dynamics. For Friends I, nothing much interesting happens: discounting causes the limiting state to be degenerate; with noise, the system may jump from one such state to the other, which does not change which states are stochastically stable.

For Friends II, as long as the number of players n is at least 4, the introduction of noise does indeed change the set of stochastically stable states: it gets rid of stars. Simulations show that pairings are by far the most prevalent states in discounted Friends II, with a star of size 3 forming when necessitated by an odd number of players. We now show that states with more than one star, or a star of size greater than 3, are not stochastically stable.

THEOREM 5. *In Friends II, with discounting, with n players, and with noise tending to zero, the stochastically stable states are those that are either unions of pairs (if n is even) or pairs plus a single star of size 3 (if n is odd).*

Sketch of Proof: Let S denote the graphs corresponding to possible limit states as in the proof of *Theorem 4*, and let $S_0 \subseteq S$ denote those graphs with no stars (perfect pairings) or with a single star of size 3. The important properties of the relation of S to S_0 are as follows. (Property *i*) If G is the result of adding a single edge to a graph in S_0 , then G contains no graph in $S \setminus S_0$. (Property *ii*) For any $G \in S$ there is a chain $G = G_1, G_2, \dots, G_k$ leading to S_0 , where each G_{j+1} may be obtained from G_j by adding an edge and then deleting two edges. Property *i* is apparent. To verify property *ii*, note that if $H \in S$ and i and j are noncentral vertices in stars of H , and they are not both in the same star of size 3, then adding the edge between i and j and removing the two edges previously incident to i and j produces a new graph in S . Iterating this procedure

starting from $H = G_1$ leads in finite time (because the number of edges decreases each time) to an element of S_0 .

We now follow the usual method for determining stochastic stability (33). Let the probability ρ of disobeying the structural dynamics Eq. 1 be very small. If ϵ (in the definition of S) is very small, then a state p with $G(p) = G \in S$ will have $G(p(t)) = G$ for all later times with high probability, until there is a disobeying move. After a single disobedience, the graph $G(p)$ will be the union of G with one extra edge. By the remark after the proof of *Theorem 4*, we see that after a disobedience, the graph will then relax to some subgraph in S . By property *i*, if $G \in S_0$ then this subgraph is again in S_0 . Thus, a single disobedience followed by relaxation back to S will never escape S_0 . Hence, the probability of jumping to $S \setminus S_0$ is of order ρ^2 , which implies that states in S_0 stay in S_0 for time at least ρ^{-2} . On the other hand, by property *ii*, from any state in $S \setminus S_0$, there is a chain of single disobediences, such that allowing the system to relax after each may with positive probability land you back in S_0 . Thus, the expected time spent in $S \setminus S_0$ before returning to S_0 is at most of order ρ . Thus, the process spends $(1 - \rho)$ portion of the time in S_0 , and sending ρ to zero, we see that only states in S_0 are stochastically stable. It is easy to see that all of these are indeed stochastically stable. QED

Reinforcement by Games of Nontrivial Strategy

So far we have only considered a baseline model of uniform reinforcement, which turned out still to have nontrivial structural behavior. Now we examine a reinforcement scheme resulting from the payoff of a nontrivial game. We will consider the case where evolution of strategy is slower than evolution of structure. Thus, we will consider the agents as divided into types, each type always playing a fixed strategy, and see what sort of interaction structure emerges. We then extend this by allowing strategic switching of types. We find that coordination of strategy occurs, though whether players coordinate on the risk-dominant or payoff-dominant strategy depends on parameters of the model such as the rate of strategic evolution. Depending on conditions of the model, the social network may or may not split up into pairs.

Rousseau's Stag Hunt. Consider a two-player version of Rousseau's *Stag Hunt* (34). The choices are either to hunt stag or to hunt rabbit (hare, in the original). It takes two persons cooperating to effectively hunt a stag, while one person acting independently can hunt a rabbit. Bagging a stag brings a greater payoff.

	Hunt Stag	Hunt Rabbit
Hunt Stag	(1, 1)	(0, .75)
Hunt Rabbit	(.75, 0)	(.75, .75)

There are two equilibria in this game: both hunt stag and both hunt rabbit. The first carries the higher payoff and is said to be *payoff dominant*; the second carries the least risk and is said to be *risk dominant* (35). In models without structural dynamics, Kandori, Mailath and Rob (36) have shown that only the risk dominant equilibrium of a two player coordination game is stochastically stable. In the presence of structural dynamics, we will describe a more optimistic conclusion.

THEOREM 6. *Suppose Stag Hunt is played by $2n$ players, with structural dynamics given by Eq. 1 and cumulative weighting dynamics (2) with no noise or discounting. Then in the limit, stag hunters always visit stag hunters and rabbit hunters visit rabbit hunters.*

Sketch of Proof: First note that no visit of a stag hunter to a rabbit hunter is ever reinforced. Thus, $w_{ij}(t) = 1$ for all t if i is a stag hunter and j is a rabbit hunter. Observing that the weights $w_{ij}(t)$ go to infinity when i and j are both stag hunters, we see

that the probability of a stag hunter visiting a rabbit hunter goes to zero.

Next, consider the subpopulation of rabbit hunters, call it A . For $i \in A$, let

$$Z(i, t) = \frac{\sum_{j \notin A} w_{ij}}{\sum_{j=1}^n w_{ij}}$$

denote the probability of visiting a given rabbit hunter visiting a stag hunter on the next turn. The expected value of $Z(i, t + 1)$ changes according to the formula

$$\mathbf{E}(Z(i, t + 1) | Z(i, t)) = Z(i, t) + t^{-1} Y(i, t),$$

where $Y(i, t)$ is the proportion of increase in expected weight w_{ij} due to $j \notin A$:

$$Y(i, t) = \frac{\sum_{j \notin A} p_{ij} + p_{ji}}{\sum_{j=1}^n p_{ij} + p_{ji}}.$$

Ignoring the terms p_{ji} in both the numerator and denominator of the above expression would lead to exactly $Z(i, t)$. The terms p_{ji} for $j \notin A$ are known to be small, while the total from the terms p_{ji} for $j \in A$ cannot be small. Consequently, $Y(i, t) < (1 - \epsilon)Z(i, t)$ for some $\epsilon > 0$, whence

$$\mathbf{E}(Z(i, t + 1) - Z(i, t) | Z(i, t)) \leq -\frac{\epsilon Z(i, t)}{t}.$$

Because the increments in $Z(i, t)$ are bounded by C/t , there are a $\lambda, \mu > 0$ for which $\exp(\lambda Z(i, t) + \mu \log t)$ is a supermartingale, which implies that $Z(i, t)$ converges to zero exponentially fast in $\log t$. QED

Introduction of a discount rate changes this outcome. Stag hunters still end up visiting stag hunters, because even discounted reinforcement beats a reinforcement of zero, but now rabbit hunters will get locked either into pairs and stars as in *Making Friends*, or into repeated visits to a single stag hunter. These limit states are all invariant under introduction of noise. When a rabbit hunter visits a stag hunter the loss to society is the 0.75 that another rabbit hunter would have profited from the visit. The model is evidently weak here, because it allows only one visit by each agent but any number of visits to each agent in a round of visits. That is, a more realistic loss would be the stag hunter's wasted time when visited by the rabbit hunter.

It should be noted that, although the stochastically stable states include those that are not optimally efficient, the optimally efficient states (those states where rabbit hunters visit rabbit hunters) will have an edge. Because of the possibility of reciprocal reinforcement, it will be easier for a rabbit hunter to switch from visiting a stag hunter to visiting a rabbit hunter than *vice versa*. Second, when the discount rate is near 1, the model behaves like the undiscounted model for a long enough time that it is very unlikely for a rabbit hunter to get locked into visiting a stag hunter in the first place. Simulations of Stag Hunting with 10 players and $d = 0.9$, seem to show that rabbit hunters "always" visit rabbit hunters. Because of both of the effects mentioned above, the system is nearly always found in an optimally efficient state, even though there are stochastically stable states that are not optimally efficient.

Coevolution of Structure and Strategy. To the previous model, we now add the possibility of an agent switching states: a stag hunter may decide to become a rabbit hunter, or a rabbit hunter may become bold and hunt stag. When this kind of strategic evolution is faster than the structural evolution, we know from studies of random encounter models that the risk dominant equilibrium of everyone hunting rabbits will be arrived at while the network is still near its initial state of uniform visitation probabilities.

Whether strategic dynamics are faster or slower than structural dynamics depends, of course, on the activity being modeled; sometimes interaction structure is externally imposed, while sometimes it is more easily modified than strategy or character. Let us suppose that the investment in re-training as a different kind of hunter is great, so between each round of visits there is only a small chance that one of the hunters will change types. Then we have seen that hunters always (with no noise or discounting) or nearly always (in discounted models) hunt with others of like type. This eliminates the risk inherent in random encounters and allows hunters to profit from switching to stag hunting after an initial period where they find another stag hunter. Slow strategic adaptation gradually converts rabbit hunters to stag hunters and the payoff dominant strategy dominates.

We describe here the results of simulations of Stag Hunting for 1,000 time steps, where with some probability q at any given time, an individual changes type to whichever type was most successful in the previous round. When $q = 0.1$, we found that in 22% of the cases all hunters ended up hunting stag, while in 78% of the cases, all hunters hunted rabbit. Thus there was perfect coordination, but usually not to the most efficient equilibrium. On the other hand, when $q = 0.01$, the majority (71%) of the cases ended in the optimal state of all hunting stag, while 29% ended up all hunting rabbit. Increasing the initial edge weights made it far less likely to reach the stag hunting equilibrium, since stag hunters took a long time to perfectly align, and without alignment, the previous round's best strategy was almost always rabbit hunting. For instance, if the initial weights were 1,000 for each visit, under 1% of the cases ended up all stag hunting, whether q was 0.1 or 0.01.

Once hunters largely cease to visit hunters of opposite type, the structural evolution within each of the two subpopulations is a version of Friends II. The resulting social structure will not be a perfect pairing, but will have each rabbit (stag) hunter visiting each other rabbit (stag) hunter, but with varying probabilities.

Conclusion

We have taken some basic steps in exploring dynamics of evolution of interaction structures and coevolution of structure and strategy. The ultimate goals are to create models that are more true to life, and to find theoretical bases for observed behaviors of systems, including prediction of selection between multiple equilibria.

The particular dynamics we use here are only examples, but it turns out that the simplest of these may deliver interesting

and surprising results. Even in baseline models where the game being played is degenerate, we find spontaneous emergence of structure from uniformity and spontaneous emergence of uniformity from structure. We find processes with extremely long transient modes, where limiting behavior is not a good guide for predicting behavior after thousands of trials.

The social interaction structures that emerge tend to separate the population into small interaction groups within which there is coordination of strategy. This separation may be complete, as in discounted Friends II, or may be only a tendency, as in the nondiscounted versions of Friends and Stag Hunting.

When we combine structure and strategy dynamics for a non-trivial game, the Stag Hunt, we find that the probable outcomes depend on the timing. Where structure is frozen in a random encounter configuration we get the expected risk-dominant equilibrium outcome. But when structure is fluid relative to strategy, structural adaptation neutralizes the risk and we get the socially efficient payoff dominant equilibrium. Varying between these extremes can give one or the other result with different probabilities, or may leave the group in a state where both strategies are used. We expect to see structure dynamics making a difference in other games as well. Indeed, we have some preliminary simulation evidence showing this to be true for a bargaining game ("split the dollar"), and for a simple coordination game.

There are many more avenues to pursue. As mentioned in Making Enemies, it would be desirable to find a model in which positive and negative reinforcement are present, but trapping does not occur. We have not modeled any interaction among three or more players. We also have yet to model any explicit interaction between strategy and structure: the choice of a partner to play with and a strategy to play against that partner need not be independent.

One could continue adding complexity so as to allow information to affect structural evolution, to include communication between players, and so forth. Our main point is this. Structural change is a common feature of the real world. A theory of strategic interaction must take account of it. There is a mathematically rich theory which develops relevant tools. We believe that explicit modeling of structural dynamics, and the interaction of structure and strategy, will generate new insights for the theory of adaptive behavior.

We wish to thank Persi Diaconis, Joel Sobel, and Glenn Ellison for bringing us together, for helpful discussions, and for greatly improving our awareness of the relevant literature. This work was supported in part by National Science Foundation Grant DMS 9803249.

- Hofbauer, J. & Sigmund, K. (1988) *The Theory of Evolution and Dynamical Systems* (Cambridge Univ. Press, Cambridge, U.K.).
- Weibull, J. (1997) *Evolutionary Game Theory* (MIT Press, Cambridge, MA).
- Feldman, M. & Thomas, E. (1987) *J. Theor. Biol.* **128**, 297–315.
- Wright, S. (1921) *Genetics* **6**, 144–161.
- Wright, S. (1945) *Ecology* **26**, 415–419.
- Jackson, M. & Watts, A. (1999) *On the Formation of Interaction Networks in Social Coordination Games*, working paper.
- Durrett, R. & Neuhauser, C. (1997) *Ann. Appl. Prob.* **7**, 10–45.
- Kang, H.-C., Krone, S. & Neuhauser, C. (1995) *Ann. Appl. Prob.* **5**, 1025–1060.
- Durrett, R. & Levin, S. (1994) *Theor. Pop. Biol.* **46**, 363–394.
- Blume, L. (1993) *Games Econ. Behav.* **5**, 387–423.
- Pollack, G. B. (1989) *Social Networks* **11**, 175–212.
- Lindgren, K. & Nordahl, M. (1994) *Physica D* **75**, 292–309.
- Anderlini, L. & Ianni, A. (1997) in *The Dynamics of Norms*, eds. Bicchieri, C., Jeffrey, R. & Skyrms, B. (Cambridge Univ. Press, Cambridge, U.K.), pp. 87–107.
- Ellison, G. (1993) *Econometrica* **61**, 1047–1071.
- Watts, D. & Strogatz, S. (1998) *Nature (London)* **393**, 440–442.
- Schelling, T. (1969) *Am. Econ. Rev. Papers Proc.* **59**, 488–493.
- Schelling, T. (1971) *J. Math. Sociol.* **1**, 143–86.
- Epstein, J. & Axtell, R. (1996) *Growing Artificial Societies* (MIT/Brookings, Cambridge, MA).
- Zajonc, R. B. (1968) *J. Personality Soc. Psychol. Monogr.* **9**, 1–28.
- Arthur, W. B. (1989) *Econ. J.* **99**, 116–131.
- Eggenberger, F. & Pólya, G. (1923) *Zeit. Angew. Math. Mech.* **3**, 279–289.
- Johnson, N. & Kotz, S. (1977) *Urn Models and Their Application* (Wiley, New York).
- Durrett, R. (1996) *Probability: Theory and Examples* (Duxbury Press, Wadsworth Publishing Company, Belmont, CA), 2nd Ed.
- Robbins, H. & Monro, S. (1951) *Ann. Math. Stat.* **22**, 400–407.
- Pemantle, R. (1990) *Ann. Probab.* **18**, 698–712.
- Benaïm, M. & Hirsch, M. (1995) *Ergodic Theory Dyn. Systems* **15**, 1005–1030.
- Talagrand, M. (1995) *IHES Publ. Math.* **81**, 73–205.
- Ehrenfest, P. & Ehrenfest, T. (1907) *Phys. Zeit.* **8**, 311–314.
- Diaconis, P. & Stroock, D. (1991) *Ann. Appl. Prob.* **1**, 39–61.
- Kac, M. (1947) *Am. Math. Monthly* **54**, 369–391.
- Davis, B. (1990) *Prob. Theor. Rel. Fields* **84**, 203–229.
- Foster, D. & Young, H. P. (1990) *Theor. Pop. Biol.* **38**, 219–232.
- Ellison, G. (2000) *Rev. Econ. Studies*, in press.
- Rousseau, J.-J. (1984) *A Discourse on Inequality*; trans. Cranston, M. (Penguin, London).
- Harsanyi, J. & Selten, R. (1988) *A General Theory of Equilibrium in Games* (MIT Press, Cambridge, MA).
- Kandori, M., Mailath, G. & Rob, R. (1993) *Econometrica* **61**, 29–56.