Stochastic Network Structure, Mobility and Efficiency

Alexander Matros
Moore School of Business
University of South Carolina

Scott Moser
Department of Government
University of Texas at Austin

April 24, 2011

Abstract

In this paper we present a simple evolutionary model of mobile agents where different 2x2 games exist at different locations. The role of information, mobility, and the payoff structure is examined for achieving global efficiency. We examine a setting where individuals are mobile and may relocate, but information about the existence and prospects of moving arrive stochastically.

1 The Model

For $t \in T$ (T countable), let $\Psi_t = (\{1, ..., M\}; E_t)$ be a finite directed network with vertex set $\{1, ..., M\}$, $(M < \infty)$ and (directed) edge set $E_t = \{i \to j \mid i, j \in \{1, ..., M\}\}$. Suppose that at each vertex $l \in \{1, ..., M\}$ there is a 2×2 coordination game G_l with the following payoff structure

$$G_l = \begin{array}{c|c|c|} & A & B \\ \hline A & a_l, a_l & 0, b_l \\ \hline B & b_l, 0 & b_l, b_l \\ \hline \end{array}$$

Figure 1: A Game at Location l

where $a_l > b_l > 0$, for any $l \in \{1, ..., M\}$. The set $\mathcal{G} = \{G_1, ..., G_M\}$

is hence a collection of M symmetric coordination games with 2 pure strategy equilibria each, a Pareto efficient equilibrium (A, A) and a risk-dominant equilibrium (B, B). This collection of games in general enough to capture various tradeoffs between efficiency and risk.[1] We call $\mathcal{G}$ the *institutional landscape* – it is the set of games possible for play by agents.

To simply the analysis, assume there are no payoff ties: $a_l \neq a_j$; $b_i \neq b_m$ and; $a_l \neq b_i$ for all $l, j, i, m \in \{1, ..., M\}$, $l \neq j$, $i \neq m$.

Any directed network whose vertex set is an institutional landscape, can be described as a *location structure* .

Definition 1 *A* **location structure** $\Psi = (V, E)$*is a directed graph with vertex set* $V = \mathcal{G}$ *and edge set* $E \subseteq \mathcal{G} \times \mathcal{G}$*, for an institutional landscape* $\mathcal{G}$

Since the vertex set of Ψ is a set of coordination games, we refer to *vertices* and *games* interchangeably. A location structure "connects" coordination games at different locations together : edges of Ψ may be thought of informational flows (so that agents at adjacent vertices may be knowledgeable of the existence, payoffs, and past behavior of agents at that vertex); or in geographic terms, so that agents at one vertex (game) are able to move to adjacent vertices (games).

We assume that in each period t, a new directed graph Ψ_t is realized. Specifically, letting $\mathcal{P}(\mathcal{G})$ be the set of all directed graphs on $\mathcal{G}$ we assume that Ψ_t is drawn according to some distribution F with full support so that $F(\Psi) > 0$ for all $\Psi \in \mathcal{P}$.[2] The stochastic generation of directed graphs in this setting may be interpreted as either (i) in each period, a graph is drawn from some probability distribution F, or (ii), in each period t, there is positive probability of an edge connecting location l and j, for any $l \neq j$.

Putting the above together, we have the following definition.

Definition 2 *A* **stochastic institutional process** *is a pair* $(F, \mathcal{G})$*, where* $\mathcal{G}$ *is a set of* M *symmetric 2x2 games and* F *is a distribution over* $\mathcal{P}(\mathcal{G})$*.*

In words, a stochastic institutional process is a way of realizing a location structure at any point in time. Call a generic network $\Psi \in \mathcal{P}(\mathcal{G})$.

[1] In section XX we show that our main results hold for a more general class of games, namely 2x2 symmetric games with two pure-strategy equilibria.

[2] In Section XX we relax this assumption.

We assume that there are $N \geq 2$ agents located at vertices $l \in \{1, \ldots, M\}$ and playing respective games. In period t, all agents at vertex l play the coordination game G_l. Furthermore, agents are mobile and can move from one game to another according to the directed network, $\Psi_t = (V_t, E_t)$. If at period t an agent is at vertex l she can move to the *neighborhood* of vertex l in the next period. The *neighborhood* of vertex l, given graph Ψ_t, is vertex l and all "nearby" vertices. A vertex j is "nearby" if there exits a directed edge $\{l \to j\} \in E_t$ from vertex l to vertex j,

$$C_l^{\Psi_t} = \{l\} \cup \{j \mid \{l \to j\} \in E_t\}.$$

Denote the neighborhood of location l in graph Ψ by C_l^{Ψ}.

In period t agents located at vertex l are randomly matched with other agents at vertex l to play coordination game G_l. If only one agent is located at some vertex, she receives a reservation utility u_r, which is less than the payoff in any pure strategy Nash equilibrium, $0 \leq u_r < \min\{b_1, ..., b_M\}$. This is a standard assumption, see for example Ely (2002), which ensures that agents strictly prefer to 'play with others' than 'leave the society'.

We assume agents are locally informed in the aggregate. Agents at vertex l at time t observes the distribution of actions taken at all games in her neighborhood: that she observes the distribution of actions of each $G_j, \in C_l^{\Psi_t}$. Based on this information she revises her current location and/or action by myopically best replying: first choosing a game in her neighborhood to play, $j \in C_l^{\Psi_t}$, and an action to play in that game. If an agent is indifferent, we assume she is equally likely to to take any action that maximizes expected payoff. This defines an agent's revision response.

This process results in a state space describing the number of agents locating at each vertex, and the distribution of action choices at each vertex. Since for a given institutional landscape, Ψ, there are M distinct 2x2 games, the state space may be written as

$$\mathcal{H} = \begin{pmatrix} n_1, n_2, \ldots, n_M \\ n_1^A, n_2^A, \ldots, n_M^A \end{pmatrix},$$

where $n_i \geq 0$ is the total number of agents locating at vertex i and $n_i \geq n_i^A \geq 0$ is the number of agents locating at vertex i, playing the efficient action A (the number of agents playing action B is $n_i^B = n_i - n_i^A$), and $n_1 + n_2 + \ldots + n_M = N$. Call a generic state h.

The process transitions from one state to another according to agent's best-reply. Given a state s and a realized location structure, Ψ, agents located at location l best respond by

1. Calculating the optimal action given the current distribution of actions for each game G_m in one's neighborhood $m \in C_l^{\Psi}$.
2. For each game in an agent's neighborhood, the corresponding expected payoff is calculated. This yields a set of expected payoffs, one for each game in one's neighborhood.
3. If this set has a unique maximum, an agent's best-reply is determined. If there are payoff ties, an agent randomly selects one location yielding maximal expected payoff, re-locates there and best-responds in that game.

Formally, a best reply is a location, l, and an action, k (in game G_l). Hence, the best-reply of agents at location l_0 given state s and location structure Ψ may be written

$$BR_{l_0}(s; \Psi, \mathcal{G}) = \underset{(l,k)}{\operatorname{argmax}}\, \pi(l,k) \text{ subject to } l \in C_{l_0}^{\Psi}$$

where

$$\pi(l,k) = \frac{1}{n_{l_0}} \left(n_l^A u_l(k,A) + (n_l - n_l^A) u_l(k,B) \right)$$

is the expected payoff of playing action $k \in \{A, B\}$ at location l.

The timing is as follows:

In period t, network structure Ψ_t is realized according to some distribution F with full support.

Agents observe the state s according to Ψ_t. For agents at location l, they observe $s_l^{\Psi} = s|C_l^{\Psi}$, where $s|V$ is the restriction of s to only vertices in V.

Agents at location l best respond to s_l^{Ψ}. (relocate, adjust strategy, both, neither)

A new network, Ψ_{t+1} is realized and repeat. This describes how the system transitions from one state to another. The state space and probability of transitions from one state to another is hence specified by $\mathcal{G}, F, N$. We write this process as $P^{\mathcal{G},F,N}$

We assume that the process starts at an arbitrary initial state, h_0: agents are randomly assigned to vertices and actions, and the system transitions according to $P^{\mathcal{G},F,N}$. Call the state of the process at time t by h_t. We are interested in both short and long-run outcomes of this process.

An example of our model, may be seen FIgure **??**.

There are two features of this modeling approach worthy of mention. First, the network structure Ψ_t describes for every period t the amount of information available to the agents: agents only have information about the (aggregate) action of agents at their respective neighborhoods. Hence, the network structure describes the flow of information and establishes the symmetry or asymmetry of agents' information. At the same time, the network structure has a geographic interpretation: it constrains agents' mobility. During a revision opportunity, agents can only 'relocate' to the neighboring locations. Hence, the network describes and limits both the information available to the agents as well as agents' mobility.[3]

2 Short Run

In this section we analyze behavior in the short run. In the short-run, agents myopically best-respond to the current state without error. We start by examining the properties of the unperturbed Markov process $P^{\mathcal{G},F,N}$. In what follows, we will make use of the following definitions. An *absorbing set* , H, is a set of states such that there is zero probability of transitioning from any state in the set to any state outside, and there is a positive probability of moving from any state in the set to any other state in the set. We take absorbing sets to be minimal w.r.t set inclusion. A state h is *absorbing* if it is a singleton absorbing set. Absorbing sets are short-run predictions: given initial conditions and the myopic-best-reply dynamic, they are states that, once arrived at, the system never leaves. A particular type of state deserves our special attention.

Definition 3 *A* **convention** *is a state such that*
(1) There is some l *with* $n_l = N$ *and* $n_i = 0$ *for* $i \neq l$ *and (2)* $n_l^A \in \{0, N\}$

A convention is a state where all agents are at the same location, taking the same action. This definition is the natural way to think of conventions in our model.

[3] Of course, we could model the flow of information with one network structure, and the mobility of agents with another network. Such separation of information and mobility could be an interesting future extension.

Our first result characterizes all short-run predictions and establishes that conventions and only conventions are short-run outcomes.

Theorem 1 *(1) All absorbing sets of $P^{\mathcal{G},F,N}$ are singleton. (2) A state is a convention if and only if it is a absorbing state.*

Proof. (1): Suppose not. There there are two states, s and s' with positive probability of transitioning from s to s' and vice versa. There are two cases to consider:

(a) All agents are located at the same vertex in states s and s' or;

(b) there are positive number of agents at at least two separate vertices.

In case (a), if all agents are at one vertex, then for s and s' to be distinct states, agents must be playing different actions in the two states. However, all agents at the same location – call it l – have the same best-reply, namely to play A if $n_l^A > \tau_l$; B if $n_l^A < \tau_l$ and; $\{A, B\}$ if $n_l^A = \tau_l$, where $\tau_l = (b_l n_l)/a_l$. Hence, generically, if all agents are at a single location, the best-reply dynamic leads to all agents playing the same action.

Case (b) implies that generically there are locations l, j and a time t such that $n_l > 1$ and $n_j > 1$.[4] Without loss of generality, let

$$\max_k \pi(l, k) < \max_k \pi(j, k)$$

so that payoff superior play is a best-response in game G_j. With positive probability, $E_{t+1} = l \to j$, so that the location structure in time $t + 1$ is a single edge from l to j, in which case the best-reply dynamic implies that $n_l = 0$ in time $t + 1$. If there are more than two states in the absorbing set, repeat this argument until all agents are at a single location, in which the arguments in case (a) apply (if $0 < n_j^A < N$ in some period $t' > t$, then again the best-reply dynamic implies that $n_j^A(N - n_j^A) = 0$ in period $t' + 1$).

(2): (only if) Obvious, as $a_l > b_l > u_r$ for all l by assumption.

(if) Let s be an absorbing state. By the argument in (1a), $n_l = N$ for some location l. Again by (1a), the best-reply dynamic then implies that $n_l^A(n_l - n_l^A) = 0$.

■

Since all conventions are absorbing states, and these are the only absorbing states, we have the following corollary.

[4] Where we ignore 'lone agents' – single agents at a location.

Corollary 1 *There are $2M$ absorbing states.*

Despite the potentially vast institutional structure, only a small number of states are short-run predictions.

3 Long Run

perturbed process and mistakes

The process described above is deterministic in nature: given an initial state of the world, we may know with certainty how the system will proceed. However, we may wish to incorporate mistakes and experimentation into the behavior of the agents, and ask what types of outcomes will stable. As such, in this section we allow agents to make mistakes and focus on predictions as the probabilty of mistakes goes to zero. We describe these predictions as long-run predictions. Long run forces may act on the behavior of the agents and "push" certain outcomes to be more likely than others, in the presence of such evolutionary forces. Accordingly, we now turn our attention to the perturbed version of the best-reply dynamic in which with probability $(1-\epsilon)$ agents best-reply to the current state and location structure and with probability ϵ agents "experiement" by choosing a location in one's neighborhood and action in that game at random. In particular, we are interested in the limiting distribution of this process as the experimentation probability tends to zero. By arguments similar to those in Young (1993), the perturbed process $P^{\mathcal{G},F,N,\varepsilon}$ is a *regular perturbation of* $P^{\mathcal{G},F,N}$, and hence it has a unique stationary distribution μ^{ε} satisfying the equation $\mu^{\varepsilon}P^{\Gamma,F,N,\varepsilon}=\mu^{\varepsilon}$. Moreover, by Theorem 4 in Young (1993), $\lim_{\varepsilon\to 0}\mu^{\varepsilon}=\mu^{0}$ exists, and μ^{0} is a stationary distribution of $P^{\Gamma,F,N}$.

example of mistakes

The following concepts are due to Freidlin and Wentzell (1984), Foster and Young (1990), and Young (1993). A state h is *stochastically stable* relative to the process $P^{\Gamma,F,N,\varepsilon}$ if $\lim_{\varepsilon\to 0}\mu^{\varepsilon}(h)>0$. Hence, long-run predictions are precisely the stochastically stable states of $P^{\Gamma,F,N,\varepsilon}$. We identify long predictions of our model with stochastically stable states because over the long run, agents might respond imperfectly (ie with mistakes) but, as usual in evolutionary models, we imagine that sub-optimal behavior is selected against, and over enough time disappears from the population.

The following theorem describes the long-run behavior of the perturbed process as the experimentation probability tends to zero.

Surprisingly, there is a unique stochastically stable state that is efficient, except in a knife-edge situation. To formalize this it will be helpful to identify the location at which the overall highest payoff is in equilibrium. To that end, order all equilibrium payoffs in descending order, $\pi_{(1)}, \pi_{(2)}, \pi_{(3)}, \ldots, \pi_{(2M)}$, so that $\pi_{(1)}$ is the highest equilibrium payoff: $\pi_{(1)} = a_{l_1} : a_{l_1} > a_l$ for $l \neq l_1$. Let l_1 be the location with the overall highest payoff. Likewise, let $\pi_{(2)}$ be the second-highest equilibrium payoff in $\mathcal{G}$: $\pi_{(2)} = c$ such that

$$c < \pi_{(1)}$$
$$c \geq \max(\bigcup(a_l \cup b_l) \setminus \pi_{(1)}).$$

And call the location at which payoff $\pi_{(2)}$ is in equilibrium, l_2. Note that if $\pi_{(2)} = b_{l_1}$, then $l_2 = l_1$.

Theorem 2
For sufficiently large N, 1. $P^{\Psi,F,N,\epsilon}$ has a unique stochastically stable state.

2a. If $b_{l_1} = \pi_{(2)}$ and (b_{l_1}, b_{l_1}) is the risk-dominant eq. payoff in game G_{l_1} then all agents locating at l_1 playing B is the stochastically stable state.

2b. If either $b_{l_1} \neq \pi_{(2)}$ or (b_{l_1}, b_{l_1}) is not the risk-dominant eq. payoff in game G_{l_1} then agents locating at l_1 and playing A is the stochastically stable state.

References

[1] Anwar (2002): "On the Co-existence of Conventions," *Journal of Economic Theory* 107, 145–155

[2] Andreas Blume and Ted Temzelides (2003): "On the Geography of Conventions," *Economic Theory* 22, 863-873.

[3] Glenn Ellison (2000): "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies* 67, 17-45.

[4] Jeffrey Ely (2002): "Local Conventions," *Advances in Theoretical Economics* 2.

[5] Foster, D. and P. Young (1990). "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology* 38, 219-32.

[6] Freidlin, M. and A. Wentzell (1984). *Random Perturbations of Dynamical Systems.* Springer-Verlag.

[7] Harsanyi, J. and R. Selten (1988): *A General Theory of Equilibrium in Games*, Cambridge, MIT Press.

[8] J. F Laslier (1997): *Tournament Solutions and Majority Voting*, Springer Verlag.

[9] Kandori, M., Mailath, G., and Rob R. (1993). "Learning, Mutation and Long-Run Equilibria in Games," *Econometrica* 61, 29-56.

[10] Nöldeke, G., and Samuelson L. (1993). "An Evolutionary Analysis of Backward and Forward Induction", *Games and Economic Behavior* 5, 425-454.

[11] Jorg Oechssler (1997): "Decentralization and the coordination problem," *Journal of Economic Behavior & Organization* 32, 119-135

[12] Young, P. H. (1993). "The Evolution of Conventions," *Econometrica* 61, 57-84.

[13] ——— (1998). *Individual Strategy and Social Structure.* Princeton University Press.