

Emergent Coordination among Competitors*

AJ Bostian[†] and David Goldbaum[‡]

Abstract

Crawford and Haller (1990) describe a repeated two-player coordination game defined by the absence of a common language. Coordination is achieved only through path dependent play relying on time consistent labels. We consider a game played by a large population similarly looking to coordinate but without the consistency in labels over time and with asymmetric coordinated payoff so that players have differing preferences regarding which coordinated structure emerges. In experiments, we link subjects together in a social network with limited ability to observe others. The complexity of the game and multiplicity of states thwarts solving for optimal play and yet the population demonstrates success in employing path dependency and the consistency of the social relationships to learn to coordinate. To capture this evolution, we model decisions with an experience-weighted attractor having recency, reinforcement, and lock-on biases. We find considerable heterogeneity in biases across individuals. Drawing on the observed biases, we conduct simulations to identify the extent to which individuals and environment determine group dynamics.

Keywords: Experiment, Simulation, Social Network, Experience Weighted Attraction, Nested Logit (JEL Codes: C73, D83, D85)

1 Introduction

This paper considers the emergence of social structures consistent with the desire to be an early adopter of a subsequently popular trend. A reward for conformity induces a population to employ social links to organize into a structured network as a means of generating coordination. In the coordinated structure, the population's leader identifies the choice for adoption. The choice disseminates to and through the followers via a network of imitations. The premium reward to acting in advance of the trend means that the number of followers and whether multiple leaders can be present in equilibrium depends on the relative strength of the reward to conformity versus the premium to leading and on the number of options faced by the population.

The investigation take place with human subjects in a behavioral lab and by simulation using virtual subjects. Whether human or virtual, players are rewarded for participating in the social phenomenon popular with their cohort and paid a premium for adopting the popular choice in advance of the phenomenon. Play

*Thanks to the University of New South Wales for the use of the ASBLab Experimental Research Laboratory and the University of Virginia for use of the Veconlab. Funding support was provided by the Australian Research Council and PureProfile through the Linkage Grant program (grant #LP0990750).

[†]School of Social Sciences and Humanities, University of Tampere, Tampere, Finland, aj@bostian.us.com

[‡]Economics Discipline Group, University of Technology Sydney, PO Box 123 Broadway, NSW 2007 Australia, david.goldbaum@uts.edu.au

takes place over a number of periods with a new set of product alternatives introduced with each period. The player has the option to adopt an alternative autonomously or through imitation of another player. Imitation delays the player's adoption but can be a mechanism for ensuring participation in an emerging phenomenon. A unique leader and population of followers is the coordinated structure that produces the cooperative outcome. All players benefit under the cooperative structure with the leader deriving greater benefit than the followers.

The experiment tests for conditions under which a population of subjects are able to use experience and recall to overcome the absence of a common language and possible competition for the leader position to identify a strategy profile that produced the desired coordination. Subjects are modeled as employing a Nested Logit model to choose between available actions. The probabilities associated with each action evolve over time as the value of each action is updated according to an Experience Weighted Attractor model of Camerer and Ho (1999). The experiments provide an opportunity to estimate the parameters of the model based on observed choices by human subjects in the lab. Simulations provide insight into how the different parameters on individual adaptive behavior affect group dynamics. Simulations using virtual subject modeled on the fitted parameters of human subjects also offer an opportunity to explore the role of individual behavior is shaping the outcome of the population dynamics.

1.1 Related literature

The players in Crawford and Haller (1990) overcome the absence of a common language by using experience and recall to identify a strategy profile that produced the desired coordination. In the current investigation, the options and labels are new to the players at each decision. Players cannot rely on their choice history to facilitate coordination but can organize to make use of stable social connections to develop paths through which decisions can decimate through the population.

The desire to adopt a popular choice means that a component of the reward is much like that which results from the strategic complementarities found in Katz and Shapiro (1985). Coordination in adoption also plays a role in the Brock and Durlauf (2001) model of conformity and the Arthur (1994) El Farol problem. Classic evidence of social influence in individual decision even in the absence of mechanical complementarities can be found in Whyte (1954), Katz and Lazarsfeld (1955) and Arndt (1967). More recent evidence involves modern technology, such as mobile phone networks as in Hill et al. (2006), and online chat rooms as in Dwyer (2007).

Early examples of the influence of social network on decisions include Schelling (1971), Schelling (1973),

Katz and Shapiro (1985), and Banerjee (1992), all of which model the bi-directional interaction between individual decisions and global behavior.¹ The architecture of a social network defines how information and influence disseminate across a population, as emphasized by Ellison (1993). Brock and Durlauf (2001) is an example for which peer effects are global, allowing tractable analysis using mean-field interactions. Cowan and Jonard (2004) document the impact of local and global connectivity on overall knowledge across a population. The leader and follower structure of the present setting emerges endogenously to serve social and individual interests.

Jackson and Wolinsky (1996) provides the setting employed in a number of models examining endogenous network formation, including Bala and Goyal (2000), Watts (2001), Jackson and Watts (2002), and Dutta and Jackson (2000), and with some modification, Kirman et al. (2007). In these models, the individual payoff function is increasing in connectivity. Players maximize connectivity subject to costs of maintaining a link. Alternatively, links can be a mean to an end not directly proportional to connectivity, as in the endogenous network formation of Galeotti and Goyal (2010). The setting shares in the asymmetry of the payoff with the core population paying more to obtain the information that is eventually shared by the entire population. Acemoglu et al. (2010) also considers the utility gain of the information that flows over the network rather than employing network connectivity as the source of utility.

The present paper determines whether a population of subjects are able to organized into the cooperative outcome represents the equilibrium of a static game as developed in Goldbaum (2014). Multiple cooperative structures exist and the asymmetry in the payoff means that the players have conflicting interests with regards to which structure emerges. The two player version of the current environment reflects the endogenous heterogeneity that can emerge in R&D and duopoly games, as in Reinganum (1985), Sadanand (1989), Hamilton and Slutsky (1990), Amir and Wooders (1998), and Tesoriere (2008). Such games may be played over two stages, but the parallel emerges in the decision regarding when the player wishes to act. Amir et al. (2010) generalizes the issue of symmetry breaking, as is the case when a leader and follower emerge. The general n player game retains the issues regarding asymmetry in outcome while introducing new strategy possibilities. It also introduces the possibility of best response cascades as in Dixit (2003) and Heal and Kunreuther (2010) as ways of refining the equilibrium set.

The relatively large population employed for the experiments means that subjects, if they play strategically, are likely to be better modeled as n two player games involving the individual playing against the population rather than one n player game, as considered in Tapiero (2014).

¹Watts (2001) and Jackson and Watts (2002) offer useful literature reviews of works on social influence.

2 Model

A population of n players face a m options in a new choice set. Call the m options alternatives. Each player chooses an action that will lead to the possible adoption of one of the alternatives. Whether an individual adopts a alternative, which one, and when is determined by the action profile of the population. Incentives reward early adoption of a popular alternative. The players lack a common language with which to refer to the alternatives. The game is played repeatedly with a new set of alternatives in each period. In this way, the game differs from Crawford and Haller (1990) for which consistency of the alternatives and their labels over time allows the players to overcome the lack of a common language through strategic path-dependent play. There is, instead, consistency in the relationships between players. Coordination, should it emerge, is achieved by exploiting the consistency of the relationships.

2.1 The game

Let $N = \{1, \dots, n\}$ be a set of players and let the $n \times n$ matrix g describe the *potential* (directed) links between players. If i can form a link to j then $g_{i,j} = 1$ and $g_{i,j} = 0$ otherwise. Let $g_{i,i} = 1$ always. Write $N^d(i; g) = \{j \in N \setminus \{i\} | g_{i,j} = 1\}$ for a set of players to which i can form a link.

Let $O_t = \{O_1, O_2, \dots, O_m\}$ be a set of alternatives. To implement the absence of a common labeling of the options, let $K = \{“A”, “B”, \dots\}$ be a set of m labels for these options and let the private one-to-one function f_i , determined by nature, map from labels to options, $f_{i,t} : K \rightarrow O_t$. Each player thus privately observes a set of labeled options $K_i = \{“A”, “B”, \dots\}$. Each player sees different labels and for every i, j pair there is a time t one-to-one correspondence, $h_{i,j,t} : K_{i,t} \rightarrow K_{j,t}$, that is unknown to the players.

Let $a_{i,t}$ denote the action of player i in period t . Players act simultaneously with each player choosing (i) one of the m options, or (ii) to link to another player. In the former case, if player i chooses option $O_{k,t}$, then assign $a_{i,t} = f_{i,t}^{-1}(O_{k,t}) = k_{i,t}$, the label by which player i refers to option $O_{k,t}$. If player i links to player j , then assign $a_{i,t} = j$. A player who chooses an option is said to *lead* while a player who links to another is said to *follow*. Thus, the set of actions for player i is $A_i = K_i \cup N \setminus \{i\}$. Write $a_t = (a_{1,t}, \dots, a_{n,t})$ for an action profile, where $a_{i,t} \in A_i$. Let $\mathcal{A}(g, m)$ be the set of all possible action profiles given g and m .

An action profile a_t induces an $n \times n$ matrix s_t describing the *actual* links between the players as determined by their actions in period t . If $a_{i,t} = j$ then let $s_{i,j,t} = 1$ and if $a_{i,t} \in K$, such that i leads, then $s_{i,i,t} = 1$. Otherwise, $s_{i,j,t} = 0$. Thus, for the matrix s , we have $s \cdot \mathbf{1} = \mathbf{1}$. Allow that j is a **predecessor of i in t** if $s_{i,j,t} = 1$ or if there is a sequence of players j_1, \dots, j_r such that $s_{i,j_1,t} = \dots = s_{j_r,j,t} = 1$. Write

$N^P(i; s_t)$ for the predecessors of i . Allow that j is a **successor of i** if $s_{j,i,t} = 1$ or if there is a sequence of players j_1, \dots, j_r such that $s_{j,j_1,t} = \dots = s_{j_r,i,t} = 1$. Write $N^S(i; s_t)$ for the successors of i in t . A **leader** is an agent who leads and has a non-empty set of successors. It is possible to *lead* without being a *leader*.

Let $N^L(s_t)$ denote the set of players who lead. If player i leads and player j is a successor of i , this makes player i player j 's leader. Note that each player i has at most one player who leads as a predecessor, that is $|N^L(s_t) \cap N^P(i; s_t)| \in \{0, 1\}$ for each i . It is possible for a successor to be without a leader. Additionally, let L_i identify the predecessor of i who is a leader.

A notion of distance is needed in order to specify payoffs.² Define the distance from a player i to her adopted option as the number of players between i and the option. Using d_i to denote player i 's distance,

$$d_i = \begin{cases} 0 & i \in N^L(s_t) \\ 1 & s_{i,j,t} = 1, j \in N^L(s_t) \\ r + 1 & s_{i,j_1,t} = \dots = s_{j_r,j,t} = 1, j \in N^L(s) \\ \infty & \text{otherwise.} \end{cases}$$

Use $D(s)$ to denote the matrix for which each row contains the ordered pair (d_i, L_i) . A second useful distance measure is the distance from a successor to a predecessor. Use $d_{i,j,t}$ to denote the distance from successor i to predecessor j measured in the number of links connecting i to j . Observe that when $L_i = j$, $d_{i,j,t} = d_{i,t}$.

Let $N^J(i; a_t)$ denote the set of players who adopt the same option as does player i (inclusive of i). Let $N^T(i; a_t)$ denote the subset of $N^J(i; a_t)$ who are of greater distance from the option than is i . Formally, $N^J(i; a_t) = \{j \in N | f_{i,t}(k_{i,t}) = f_{j,t}(k_{j,t})\}$ and $N^T(i; a_t) = \{j \in N^J(i; a_t) | d_{j,t} > d_{i,t}\}$. Let $\mu_{i,t}^J = |N^J(i; a_t)|$ and $\mu_{i,t}^T = |N^T(i; a_t)|$.

The payoff for player i is

$$\pi_{i,t} = \pi(i; s_t) = a_J(\mu_{i,t}^J - 1) + a_T \mu_{i,t}^T \quad (1)$$

with coefficients $a_T \geq 0$ and $a_J \geq 0$. The first element of the payoff rewards conformity. The second element rewards early adoption based on the number of subsequent adopters of the same alternative. distance advantage a player has over other players. Let $\Pi(a_t) = (\pi_{1,t}, \dots, \pi_{n,t})'$.

The static version of the developed game has a Nash equilibrium identified by the s , reflecting that the game is played in the strategic use of links rather than an attempt to strategically choose from the available

²The notion of time and distance are interchangeable when adoption disseminates at a rate of one unit of time per link.

alternatives. For use in the following propositions, let

$$\mu^* = \frac{1}{1 - (m-1)\frac{a_J}{a_T}} \leq 1. \quad (2)$$

The structures identified in Goldbaum (2014) as the Nash equilibria to the static game will simply be referred to as the cooperative structures of this dynamic game.

Proposition 1. *For $(m-1)a_J \geq a_T$ and for every $i \in N$, s is cooperative structure if i is a leader with a successor population $N^S(i; s) = N \setminus i$ and where, for each $j \in N^S(i; s)$, $d_{j,i}$ is the distance attained by minimizing the number of links between j and i .*

Proof. see Goldbaum (2014) Propositions 3. □

Proposition 2. *For $(m-1)a_J < a_T$ and for every $i \in N$, s is a cooperative structure if i is a leader with a successor population of size such that $(|N^S(i; s)| - \mu^*)^2 < 1$ and where, for each $j \in N^S(i; s)$, $d_{j,i}$ is the distance attained by minimizing the number of links between j and i and $d_{j,i} \leq d_{j',i}$ for all $j' \notin N^S(i; s)$.*

Proof. see Goldbaum (2014) Propositions 5. □

By Proposition 1, for a large number of options or a sufficiently large relative payoff to conformity, the cooperative structure has everyone within the population organizing around a single leader in order to maximize expected payoff in a single period given the actions of the other players. For every player, there is an equilibrium structure in which she is the leader. This structure offers the highest conformity payoff attainable. For the leader, it offers the highest possible payoff attainable to any individual from any feasible structure and for the followers, the highest possible autonomous payoff given the leader.

Under the conditions of Proposition 2, those who would otherwise be most distant from the leader find that leading offers a higher expected payoff than following. The relatively high payoff to early adoption induces them to gamble that they will independently choose the same alternative as the leader and thereby enjoy the same payout as the leader.

The determinants for the equilibrium size of the successor population are the ratio a_J/a_T and m . It is independent of the topology of g or the realization of s . For the dynamic game played by the experiment subjects, the structures described in Propositions 1 and 2 need not be equilibria. For this reason, they will be referred to as the “cooperative structures.”

2.2 Individual information and strategy estimation

Participants are players in a complex dynamic game for which the equilibrium strategy is not obvious. Strategies that increase the likelihood of emerging as the leader conflict with the strategies that help produce the cooperative structures. It is most likely that optimal strategic behavior for player i involves a degree of path dependence that revolves around what can be inferred about s_t from the actions and payoffs observed for the $N^d(i; g)$ population. At the beginning of each period, the player's information set includes the player's personally observable outcomes of all prior rounds. Let $X_{i,t} = \{x_{i,t-1}, x_{i,t-2}, \dots, x_1\}$ where $x_{i,t}$ contains information observed by player i at the end of period t . This includes the total popularity of each alternative as it evolved during the period. Players also observe alternative choice and timing of each contact's adoption within the period.

From the information in $x_{i,t}$, players can calculate their payoff for period t . They also estimate the counterfactual values of the actions not taken in that period. They cannot compute counterfactuals with certainty because the players do not know the entire structure of s_t . What they can do, for example, is take the actual adoptions as given and compute the payoff had they adoption according to each of the counterfactual actions.

The estimated empirical model uses a nested logit to assign probability to each action using path dependent performance measures, updated according to the experience weighted attractor (EWA). In period t , for $\theta_{i,t} \in [0, 1]$ and $\sum_j w_{i,t}^j = 1$, let

$$\begin{aligned}\theta_{i,t} &= \Pr(i \text{ leads in } t) \\ (1 - \theta_{i,t})w_{i,t}^j &= \Pr(i \text{ follows } j \text{ in } t).\end{aligned}$$

For agent i , if j is in her contact list, then $w_{i,t}^j \in [0, 1]$, otherwise $w_{i,t}^j = 0$.

The performance associated with each strategy option is tracked and updated according to the experience weighted attractor (EWA) of Camerer and Ho (1999). That is, there is a performance measure for agent i associated with leading,

$$\begin{aligned}A_{i,t+1}^0 &= \frac{\phi_l N_{t,l} A_{i,t}^0 + (\delta_l + (1 - \delta_l)\mathbf{1}(x_t = l))\pi_{i,l,t}}{N_{t+1,l}} \\ N_{t+1,l} &= 1 + \phi_l(1 - \kappa_l)N_{t,l}\end{aligned}\tag{3}$$

and for following agent j ,

$$\begin{aligned} A_{i,t+1}^j &= \frac{\phi_f N_{t,f} A_{i,t}^j + (\delta_f + (1 - \delta_f) \mathbf{1}(x_t = d)) \pi_{i,j,t}}{N_{t+1,f}} \\ N_{t+1,f} &= 1 + \phi_f (1 - \kappa_f) N_{t,f}. \end{aligned} \quad (4)$$

Relative performance determines the probability of adoption of an action according to the nested logit model.

For performance measures $A_{i,t}^0, A_{i,t}^1, \dots, A_{i,t}^d$,

$$\theta_{i,t} = \frac{\exp(\mu A_{i,t}^0)}{\exp(\mu A_{i,t}^0) + \left(\sum_{j \in D} \exp(\mu A_{i,t}^j / \lambda) \right)^\lambda} \quad (5)$$

$$(1 - \theta_{i,t}) w_{i,t}^j = \frac{\exp(\mu A_{i,t}^j / \lambda) \left(\sum_{j \in D} \exp(\mu A_{i,t}^j / \lambda) \right)^{\lambda-1}}{\exp(\mu A_{i,t}^0) + \left(\sum_{j \in D} \exp(\mu A_{i,t}^j / \lambda) \right)^\lambda}. \quad (6)$$

The EWA captures the players changing perception of attractiveness of each possible action. The parameters of the EWA reflect the extent to which participants update old information with new and how much credence they give to the counterfactuals. Decreasing ϕ from one towards zero shifts weight towards more recent performance realizations, producing “recency” in the weighting of past performance. Decreasing δ from one towards zero down-weights the counterfactual payoffs so that at the extreme $\delta = 0$ the agent engages in reinforcement learning. For $\kappa \rightarrow 1$, persistent performance differences accumulate over time to drive probabilities towards the extremes, known as “lock-in”.

Correlation between the choices in a multinomial logit environment requires estimation with a nested logit. The Generalized Extreme Value distribution upon which the nested logit is based allows for correlation within the “nested” options and is useful for decomposing choice options that are separated by an implicit sequential ordering of decisions in a tree structure or by some other grouping mechanism.

The parameter $\mu \geq 0$ identifies how sensitive the players are to differences in the performance measures with $\mu = 0$ indicating indifferent between the actions.

The estimate of $0 \leq \lambda \leq 1$ identifies the extent to which the player compartmentalizes the decision about whether to lead or follow from the decision about who to imitate when following. For $\lambda = 1$, the agent treats each options as independent. With $A_0 = \dots = A_d$, then $\theta_i = (1 - \theta_i) w_i^j = 1/(d + 1)$. If one option exhibits higher performance than the others, it draws probability weight away from the others options equally. For $\lambda \rightarrow 0$, the agent clusters all of the following options. With $A_0 = \dots = A_d$ the two options in the parent decision regarding whether to lead or follow are given equal weight so that $\theta_i = 1/2$.

Each option in the dependent decision about who to imitate gets equal weight within the follow option so that $w_i^j = 1/d$ for each $j \in N^d(i)$. If one contact outperforms the others, it primarily absorbs probability weight from the other contact. It also absorbs probability weight from the lead option, but only to the extent that the aggregate performance associated with following increases with the additional probability weight on the better performing contact. Nesting preserves the Independence of Irrelevant Alternatives, the condition under which for possible actions, a_i and a'_i , $\Pr(a_i)/\Pr(a'_i)$ does not change when an irrelevant option is added as a possible action.

3 The in-lab experience

Experiments were conducted at the ASBLab Experimental Research Laboratory of the University of New South Wales and at the Veconlab of the University of Virginia. Subjects were recruited using the ORSEE recruitment system.³

An experiment consists of 36 anonymous subjects seated at computer monitors. Subjects are informed that they will be playing a game for 60 periods randomly grouped with 17 other participants. In each period the group will face a list of six alternatives of which they can adopt one option in the period. The labels on the different alternatives are randomly assigned for each member of the group so that Alternative 1 may be labeled Choice 2 or player A and Choice 3 for player B. The labels are reassigned in each period. The subject will be able to observe the decisions of d other participants in their group where d is either 3 or 4, depending on the treatment. The observation list is randomly generated and changes every 20 periods. The participants on the observation list will be labeled “Person A”, “Person B”, etc.

The subjects are also informed about how decisions will be made. Each period consists of 7 subperiods. Before a period begins, each subject will make two decisions, they will select (i) one of the six choices and (ii) identify one of the members of their observation list to follow. They also identify which decision, (i) or (ii), they wish to implement for the period. Figure 1 shows an example of the decision screen faced by each subject at the start of each period.

Once everyone in a group has submitted their strategy for the period, the subperiods will begin ticking down. If a subject chooses directly, the choice will be implemented in subperiod 1. If the subject chooses to imitate a contact, they will implement the choice in the subperiod after the contact’s choice is implemented. A representation of this is presented in Figure 2. This screen is not shown to the subjects.

³See Greiner (2004)

Experiment (ID 1)
Player ID: 1

60 periods total

The list of d other players whom the subject can imitate.

3x20-period shuffles

strategy adoption area

Round 1

Player	Choice	Subperiod
A	n/a	n/a
B	n/a	n/a
C	n/a	n/a
You	n/a	n/a

You have been assigned a new observation list. This new list will be in effect for the next 20 rounds.
 Reminder: Your observation list will change every 20 rounds.

Choice	Count
1	0
2	0
3	0
4	0
5	0
6	0

Your choice list has been re-shuffled for this round.
 Reminder: Your choices are randomly shuffled at the start of each round, so that the choices reflect different alternatives. When necessary, the program will automatically translate between your choices and the underlying alternatives.

Please enter both of the following decisions:

Decision 1. A choice to make on your own: Choice 1 ▾

Decision 2. A player to follow: Player A ▾

Now, please indicate which of the above decisions you would prefer to use as your **strategy** in this round.

Decision 1 (Choose on your own.)

Decision 2 (Follow another player.)

Figure 1: Screen shot of decision page

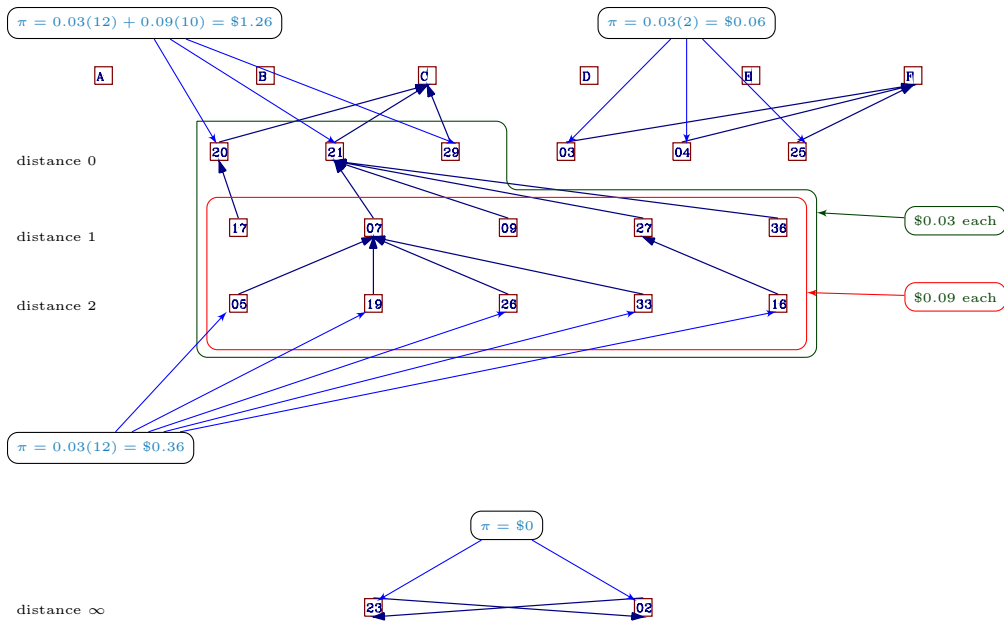


Figure 2: Behind the scenes

The subjects are also informed of the earnings formula. They told the values of a_J and a_T and that they will receive earnings in each period,

$$\begin{aligned} \text{earnings} &= a_J \times \text{Number of People in your Group Who Selected the Same Alternative as You} \\ &+ a_T \times \text{Number of People in your Group Who Selected Your Alternative after you Did.} \end{aligned}$$

At the end of each period, after the last subperiod, the subjects are shown information about the popularity of each option, the alternative implemented by each of their contacts, their own earnings for the period and a calculation of the earnings they would have received for each of the alternative actions available to them. Figure 3 is an example of the earnings screen seen at the end of each round. Across the top of the figure are the six choices, labeled using capital letters. Individual agents, labeled using numbers, appear in rows below the choices based on the distance from the adopted option.

After completing the game, the subjects provided basic demographic information and answer questions allowing scoring on the “Big 5” personality factors, cognitive reflection test, and for resistance to peer influence.

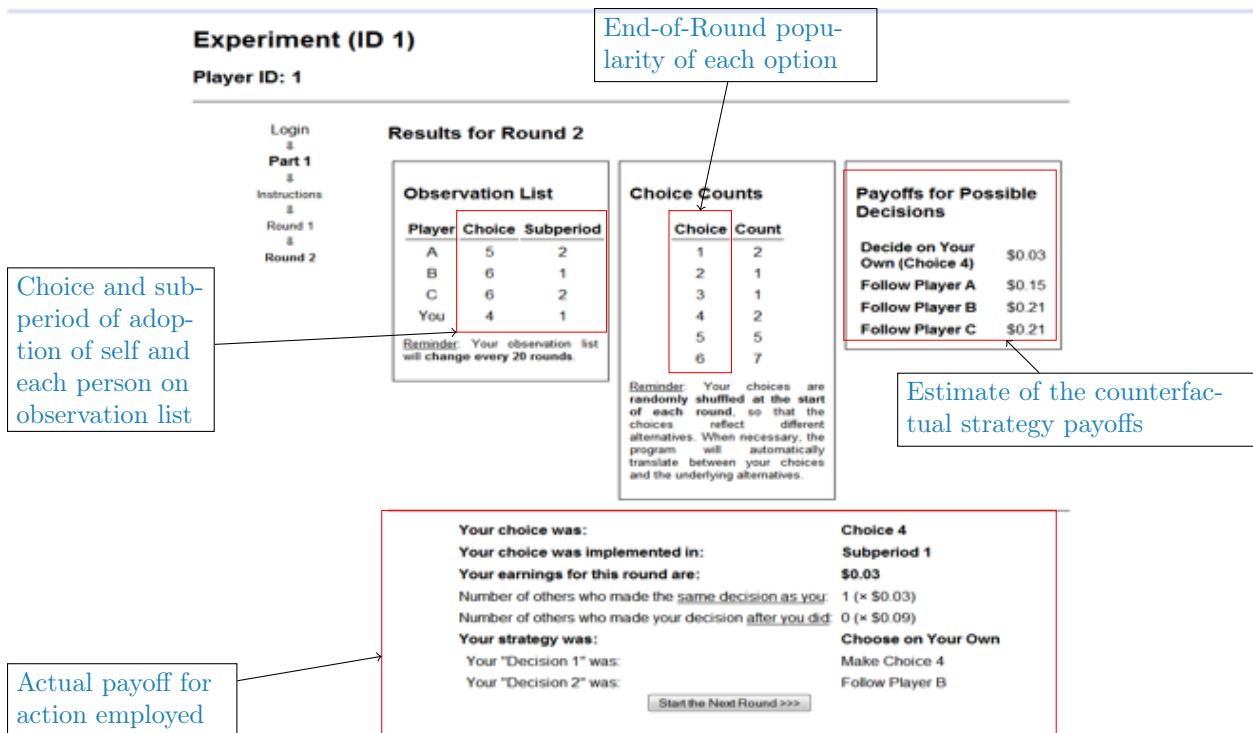


Figure 3: Screen shot of results page

Table 1: Treatments. Each session involves 36 subjects randomly allocated into two groups of 18 players each. Subject play the game for a total of 60 round made up of three shuffles of 20 periods each. A subject plays with the same randomly allocated group (and contact list) during the 20 period of a shuffle.

Session ID	Population	Rounds	Location	# Contacts	a_J	a_T	Ratio
1	2×18	3×20	AU	3	0.03	0.09	3:1
9	2×18	3×20	AU	4	0.05	0.15	3:1
10	2×18	3×20	AU	4	0.05	0.15	3:1
17	2×18	3×20	AU	3	0.03	0.09	3:1
18	2×18	3×20	USA	4	0.05	0.30	6:1
19	2×18	3×20	USA	4	0.05	0.30	6:1
25	2×18	3×20	USA	3	0.05	0.30	6:1
32	2×18	3×20	AU	3	0.05	0.30	6:1
33	2×18	3×20	USA	4	0.05	0.15	3:1
34	2×18	3×20	USA	4	0.03	0.09	3:1

A total of ten sessions were held. As reported in Table 1 the two manipulation involve altering the rates of compensation for popularity and early adoption and altering the number of members on a subject’s contact list.

4 Experiment results

We report results regarding the influences, internal and external, on whether a group forms into a cooperative structure. We start with an anecdotal narrative and simple frequency measures of success. These will provide a foundation from which to understand the estimates of the model capturing dynamic adjustment in player strategy and the role played by individual decisions and environmental factors in driving organizational outcomes.

4.1 Evolution

Of the 360 participants across the 10 different sessions, 319 lead in the first period. By the end of the first shuffle (period 20), only 177 of the participants lead. The path from an initially unstructured, unorganized group of individuals predisposed to independent adoption of the available alternatives to a structured population able to coordinate on a single alternative involves a number of steps that can be seen in the populations evolution in the experiments. It starts with some of the players choosing to follow rather than lead. If one player, say player i , can attract a persistent set of successors, $N^S(i; s_t)$, even if small, this provides others in the population the opportunity to observe the high payoffs offered by becoming a member of $N^S(i; s_{t+1})$. For $(m - 1)a_J > a_T$, the larger the successor population, the greater its visibility to the remaining unaffiliated

d \ ratio	6:1	3:1	
3	0.167 (12)	0.25 (12)	0.208 (24)
4	0.25 (12)	0.708 (24)	0.556 (36)
	0.208 (24)	0.556 (36)	0.417 (60)

Table 2: Probability of groups forming the cooperative structure. For the 3:1 payoff ratio, the coordinating structure is present if the dominate leader has 13 or more followers. For the 6:1 payoff ratio, the coordinating structure is present for 5 or more followers. Number of groups reported in parenthesis.

population and the greater the benefit to joining.

In those groups where the coordinated structure fails to form, the process typically breaks down when early followers fail to remain as successors while $N^S(i; s_t)$ is small.

4.2 Simple frequencies

For $(m-1)a_J \geq a_T$ (Proposition 1), the group is identified as having successfully coordinated if, in the final period, there is a single dominate leader with thirteen or more followers.⁴ Sessions with a payoff ratio of 6:1 for leading have $(m-1)a_J < a_T$ and $\mu^* = 6$. Success is achieved if the structure in the final period consists a dominant leader with five or more followers.

As reported in Table 2, while the coordinated structure emerges in all scenarios, the frequency of emergence depends on the treatments. The coordinated structure emerges with greater frequency when the premium earned by the leader is less extreme. Increasing d increases the number of other players each participant is able to directly observe and reduces the overall distance information needs to travel to reach everyone in the population.

Somewhat odd is the fact that it is the second shuffle out of three that enjoys the greatest success in generating the coordinated structure.

4.3 Individual play, pooled model estimate

Individual play is estimated to evolve according to the EWA and nested logit model described in equations 3 through 6. Within a session, the periods $t = 1, 21, 41$ are each the first period of a new shuffle. With A_1^d , A_{21}^d , A_{41}^d set to zero, the values of A_1^0 , A_{21}^0 , A_{41}^0 reflect the player's predisposition towards leading, which may change based on the subjects prior experience. In the baseline model, the values of A_1^0 , A_{21}^0 , and A_{41}^0 are free parameters chosen in the estimation procedure. The alternative model allows each participant and individual bias as determined by their individual characteristics and personal experiences during the game.

⁴The standard is subjective based on analysis of group behavior. Increasing the threshold to fourteen and the frequency of success is

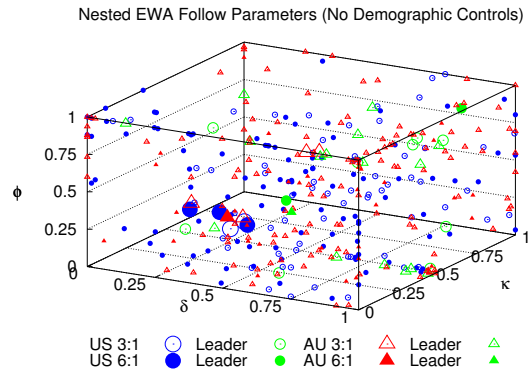

Table 3: Pooled: EWA and nested logit parameter estimates. Without characteristics, A^0 is estimated as a free parameter of the model reported as \hat{A}^0 . With characteristics, A^0 is estimated as a linear function of the individual characteristics and experiences. \hat{A}^0 is the population average of the fitted value.

	Without A_0 Characteristics			With A_0 Characteristics		
	Global	3:1	6:1	Global	3:1	6:1
ϕ_l	0.728	0.732	0.731	0.806	0.801	0.830
δ_l	0.070	0.086	0.016	0.067	0.103	0.000
κ_l	0.541	0.501	0.712	0.801	0.777	0.800
ϕ_f	0.467	0.468	0.438	0.448	0.432	0.442
δ_f	0.491	0.534	0.433	0.495	0.536	0.443
κ_f	0.000	0.000	0.000	1.000	1.000	1.000
λ	0.431	0.440	0.373	0.465	0.478	0.433
μ	3.007	3.359	2.174	1.766	2.033	1.399
\hat{A}_1^0	0.794	0.768	0.857	1.252	1.181	1.341
\hat{A}_{21}^0	0.573	0.444	0.911	1.113	0.827	1.737
\hat{A}_{41}^0	0.419	0.238	0.932	0.850	0.410	1.871
N	360	216	144	360	216	144

Table 3 reports the estimated coefficients of the EWA and nested logit model (EWANL) for both the baseline model with A_1^0 , A_{21}^0 , and A_{41}^0 estimated as free parameters and the alternate model based on individual characteristics. The reported \hat{A}_{i,t_0}^0 are the average fitted values for A_{i,t_0}^0 of the indicated population. The variables and associated coefficients for the determinants of individual A_{i,t_0}^0 are reported in Table 4.

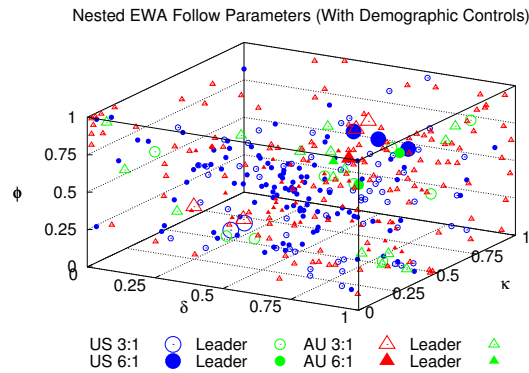

Figure 4 displays the scatterplot of the parameter estimates produced by fitting the EWANL model to individuals. Figures 4a and 4b display the lead and follow parameters respectively for the baseline model with free \hat{A}_{i,t_0}^0 parameters. Figures 4c and 4d display the EWA parameters when the individual model is estimated based on personal experiences in generating \hat{A}_{i,t_0}^0 . The variety of parameter values points to the diversity in how individuals manage the measures of performance. The green points indicate the coefficients of leaders. The distribution over the possible parameter space suggests that leadership was not the product of a player's strategy alone. The uniformity in the distribution of individuals across treatments points to the environment, rather than individuals involved, as the explanation for the different frequency of success in coordination.

Estimating the model with personal characteristics increases the concentration in the parameters, suggesting that accounting for an individual characteristics and experiences explains some of the heterogeneity in behavior.



(a) Individual leader EWA coefficient of in baseline model with free \hat{A}_{i,t_0}^0 parameters

(b) Individual follower EWA coefficient of in baseline model with free \hat{A}_{i,t_0}^0 parameters



(c) Individual leader EWA coefficient of in alternate model with \hat{A}_{i,t_0}^0 estimated using personal characteristics

(d) Individual follower EWA coefficient of in alternate model with \hat{A}_{i,t_0}^0 estimated using personal characteristics

Figure 4: Individual behavior: A scatterplot of the parameters of the EWA from estimating the EWANL model on individuals.

Table 4: Pooled: Propensity to lead.

	A_0 Characteristics	Global	3:1	6:1	AU	US	Avg.	Std Dev
Demo	age	-0.032	-0.029	-0.033	-0.033	-0.030	21.691	2.990
	female	-0.272	-0.389	0.033	-0.281	-0.288	0.425	0.495
Psync	extraversion	-0.019	-0.001	-0.066	0.092	-0.103	3.191	0.724
	agreeableness	-0.068	0.028	-0.236	-0.044	-0.007	3.631	0.621
	conscientiousness	0.046	0.039	0.044	0.004	0.093	3.439	0.636
	neuroticism	-0.036	0.049	-0.225	-0.019	-0.035	2.704	0.752
	openness	-0.060	-0.103	0.022	-0.150	0.004	3.399	0.561
	cognitive reflection	0.056	0.083	-0.053	0.053	0.091	1.631	1.117
	resistance to peer influence	0.149	0.162	0.183	0.052	0.280	2.696	0.361
Exp*	total lead decisions (out of 90)	-0.003	-0.016	0.003	-0.002	-0.003	42.1	4.74
	individual times lead (out of 5)	0.490	0.411	0.601	0.546	0.431	2.34	1.12
	avg. distance from leader	-0.173	-0.267	0.165	-0.103	-0.152	1.83	1.15
	avg. profit	-0.214	-0.479	1.290	-0.501	0.275	0.72	0.70
	distance x profit	0.230	0.249	-0.629	0.433	-0.094	1.35	1.32
N		360	216	144	180	180		

*Last 5 periods of prior shuffle

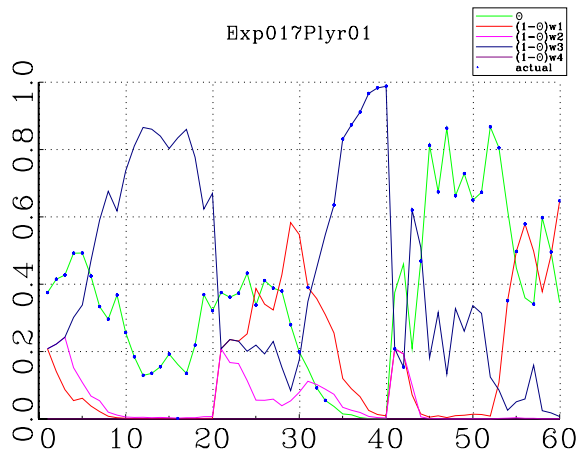
4.4 Analysis

4.4.1 Dynamic behavior

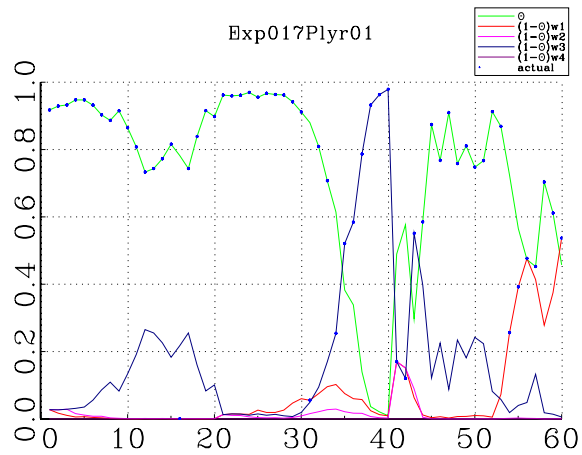
Consider the behavior of player 1 in session 17 who, over his three different shuffles, experienced both success and failure of group coordination. Session 17 has a payoff ratio of 3:1 and $d = 3$ and thus is has settings producing a 25% chance of a group organizing into the coordinated structure. Player 1's first grouping, from period $t = 1, \dots, 20$, fails to develop the coordination structure. Player 21 finishes the shuffle with four successors, but player 18 is one of a number able to attract a consistent following that lasted for a number of periods.⁵ Player 18 peaks with four successors in $t = 9$. Player 1's second grouping, $t = 21, \dots, 40$, produces the coordinated structure with player 25 ending the shuffle with fourteen followers. Player 1's third grouping, $t = 41, \dots, 60$, is unique among all sessions in the frequency and number of players involved in self-referencing loops. Player 18 is an early leader with thirteen successors in period $t = 44$ but only player 31 is directly linked to player 18. The remaining twelve successors are linked to 18 through player 31. In period $t = 45$ player 31 leads and in period $t = 48$, players 18 and 31 imitate each other, forming a self-referencing loop that includes fourteen players. A number of the periods that follow also include large self-referencing loops, disrupting the emergence of consistent, reliable personal links. The shuffle ends with player 25 having twelve successors (with only player 33 directly linked to 25).

Figure 5 depicts the time-series of fitted probabilities, $\theta_{i,t}$ and $(1 - \theta_{i,t})w_{i,j,t}$, for payer 1 of session 17.


⁵Within each session, the 36 players have identification numbers $1, \dots, 36$ that they retain while placed into groups of 18 players.



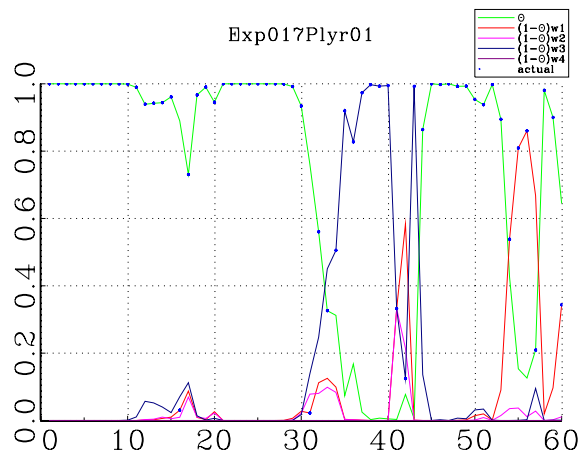
(a) Full information EWA and unbiased



(b) Full information EWA with individual bias



(c) Global estimated EWA with individual bias



(d) Individually estimated EWA with individual bias

Figure 5: Player 1 in Session 17: Fitted values for $\Pr(\text{lead}) = \theta_{i,t}$ and $\Pr(\text{follow } d) = (1 - \theta_{i,t})w_{i,j,t}$. Initial biases toward leading are $\hat{A}_{1,1}^0 = 1.65$, $\hat{A}_{1,21}^0 = 2.11$, and $\hat{A}_{1,41}^0 = 0.27$. The global parameters in frame (5c) are as reported in Table 3 with characteristics. Player 1's individual EWA parameters used in frame (5d) are $\phi_l = 0.729$, $\delta_l = 1$, $\kappa_l = 1$, $\phi_f = 0.713$, $\delta_f = 0.730$, $\kappa_f = 0.516$, $\mu = 6.167$, $\lambda = 0.370$, $A_{1,1}^0 = 16.33$, $A_{1,21}^0 = 9.78$, and $A_{1,41}^0 = -0.863$.

The dots identifies the series of the actions chosen by the participant. The probability values reflect player 17's initial biases and subsequent adaptation to the events he is able to observe around him. Each frame of Figure 5 is the same underlying events and decisions, but generated based on a different set of EWA parameters. Thus, each frame tells a different narrative of player 1's behavior.

For 5a, all of the EWA parameters are set to one, $A_{i,1}^0 = A_{i,21}^0 = A_{i,41}^0 = 0$, $\mu = 1.766$, and $\lambda = 0.467$. The latter two are the estimated values for the global model with characteristics. These are the values that a player would employ in order to make full use of all the data generated without decay and without an initial bias towards leading. With these initial conditions, the value of θ_1 is approximately 0.4 while each of the three contacts have a probability of 0.2 each of being selected. For a full information player in player 1's position, the probability of leading soon falls to around 0.2 as the counterfactual option to follow contact #3 (player 18) pays relatively well. Player 18's success is short-lived but no good alternative emerge as no single player attracts a consistent following. Persistent leading by player 1 and others in the group means that meaningful opportunities to coordinate pass without being exploited.

Although shuffle two starts much like the first, here players do respond to the apparent opportunities presented by the success of player 25 in attracting followers. This solidifies the emergent structure and reinforcing the advantage to joining the hierarchy under player 25. Player 1 does not have a direct link to player 25 but player 9 (contact #3) does. In this occasion, player 1 responds to the apparently successful strategy of following. Player 1 first follows player 2 (contact #1) in period 31, switches back to leading, and finally permanently switches to follow player 9 in period 34, largely responding to apparent opportunities suggested by the full information evaluation of the options.

The third shuffle starts differently than the previous two. Player 1, like many other participants, start the shuffle following and a dominant leader quickly emerges. Acting contrary to the signaled opportunities, player 1 switches to lead in $t = 44$ but is fortunate to choose the same alternative as does the leader. Player continues to be lucky in choosing the option chosen by the dominant leader when there is one, keeping the performance measure for leading high. Leading also helps player 1 to avoid the self-referencing loops when they occur. The increase in probability associated with contact #1 towards the end of the shuffle reflects the emergence of player 25 as a leader with a large number of successors in the aftermath of the series of self-referencing loops. Player 1's shortest distance to player 25 is three links, via players 9 (contact #1) and 33.

The parameters used to generate the series in frame 5b, are the same as in frame 5a except that $\hat{A}_{i,t_0}^0 = (1.65, 2.11, 0.27)$ for $t_0 = 1, 21, 41$, from the estimates of the global model based on player 1's personal

characteristics and experiences.

While contact #3 offered the best performance when measured without bias, the performance remains weak and is insufficient to overcome player 1's bias towards leading. Including this bias helps to explain the persistent decision to lead throughout the first shuffle and for most of the second shuffle. In the second shuffle, player 1's bias towards leading is even stronger but this time the payoff gain from following is sufficiently large to overcome the bias. Including the bias, the θ_1 is still high, but declining, when player 1 switches to follow. The smaller bias towards leading is barely discernible in the probabilities at the start of the third shuffle.

The series displayed in 5c are generated from the EWA, λ , and μ , from the global model with individual characteristics included to estimate \hat{A}_{i,t_0}^0 . Two distinguishing features are produced by $\delta_l = 0.067$ and $\phi_l = 0.806$. Approaching the extreme of reinforcement learning, the strong relative performance of contact #3 in the first shuffle does not even register due to the heavy discounting of the counterfactual to leading. The benefits of following only become apparent after the first sample of the strategy, as apparent by the jump in the probability of imitating contact #3 to $(1 - \theta_1)w_{1,9}$ in $t = 35$.

Having $\phi_l < 1$ means that the strong early bias towards leading dissipates as the approach fails to offer commensurately high payoffs, producing the observed decline in θ_1 from its high starting value in the first and second shuffles.

The series displayed in 5d are generated from the EWA, λ , and μ , from the estimate of the player 1 specific model with free \hat{A}_{i,t_0}^0 . When compared to the global parameter estimates, the relatively large value of $\mu = 6.167$ means that the probabilities will tend to extremes for even small differentials in the performance measure. This persistence in leading is reflected in the substantial value of $A_{1,1}^0 = 16.33$ but is still constrained to allow player 1 to follow in $t = 16$. The median value for $A_{1,1}^0$ for session 17 is 1.57. For the second shuffle, the bias remains large so to keep θ_1 near 1 through the first 10 periods of the shuffle before it falls quickly just before player 1 switches to follow contact #3. For the third shuffle, the bias favors following, which is what player 1 starts out doing. Player 1's switch to, and then persistence in, leading is the consequence of lucky alternative selection and the atypical collapse of the cooperative structure.

Figure 6 plots the fitted values of θ_i as estimated by the global model as realized in the third shuffle. The trends reflect the experiences of the individuals as s_t evolves around them. The frames, drawn from four different environments, capture the different behavior norms that arise across the session. The top row has the payoff ratio of 3:1 and the left column with $d = 4$. The players in the third shuffle of session 9 quickly fall in behind a single emergent leader. The spikes in the θ_i of individual followers indicate that the follower's

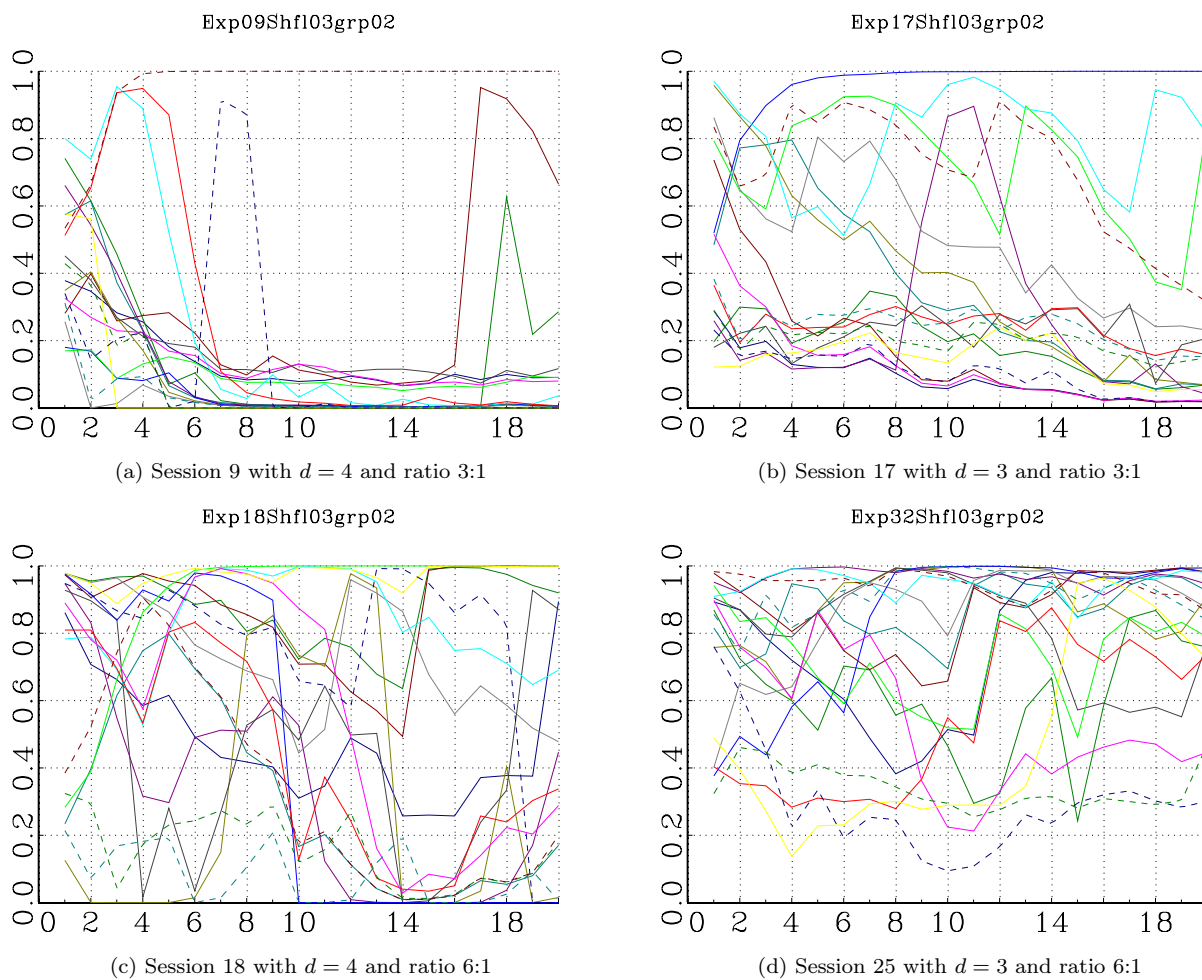


Figure 6: Examples of the evolution of θ_i from the shuffle 3 (and group 2) of the session the indicated session. In frame (a) a leader emerges early and the population quickly follows. In frame (b) a leader emerges and the population slowly learns to follow. In frame (c) a leader emerges according to the cooperative solution for a high payoff ratio but the population does not learn the lead or follow role that is appropriate to the individual within the 20 periods as a group. In frame (d) no leader emerges and no one learns to follow.

counterfactual choice coincided with that of the leader but that measure of performance quickly dissipates. In the third shuffle of session 17, the players are much slower to fall in behind the single leader despite the leader’s early emergence. This is attributable to the decrease in individual visibility and the greater distance information must travel with $d = 3$ rather than 4.

Recall that for a payoff ratio of 6:1 the cooperative structure has a single leader with six followers and that the remainder of the population should individual lead. Behavior consistent with this structure would thus produce twelve players with θ_i near one and six players with θ_i near zero. The group the third shuffle of session 18 are successful in producing a single dominant leader who has, in the final period, seven followers, but an examination of the Frame 6c reveals that the distribution of θ_i fails to conform. This is a consequence of the estimated behavior rather than the experiences of the players which, using the baseline parameters, would have produced appropriate values of θ_i at the extremes. In the third shuffle of session 25, no leader emerges able to attract more than two or three followers for more than a couple of periods. This is characteristic of those sessions where no dominant leader emerges. As seen in Frame 6d the individual θ_i values do not evolve to support persistent following, consistent with player behavior. The outcome is interesting in how the failure of a leader to emerge is self-reinforcing. For a leader to emerge requires a dedicated following, even if small in number, in order to attract more followers, reinforcing the decision to follow.

5 Simulations

Simulations allow for a comparative static-type investigation into the impact of the different environmental parameters as well as the different subject behavior as captured by the parameters of the EWA and nested logit. The majority of the simulations in this section are populated by *ex ante* homogeneous subjects. This means that the subjects all employ the same probabilistic nested logit model to determine their action for the period and EWA model with which to update performance measures based on experiences. *Ex post*, the subject will differ in the probabilities associated with each action as performance measures are updated based on personal experiences. For the simulations, each iteration involves a single “session” of 20 rounds played by a population of 18 players. Each treatment is simulated for 1000 iterations. The frequency with which the coordinated structure emerges according to the same standard applied to the human experiments is recorded. Each treatment starts with the same seed value for the random number generator.

Table 5 reports the baseline parameters of the simulation. A random process determines the contact list

Table 5: Parameters of the Base Simulation

Population	$n = 18$	Participation	$(m - 1)a_J > a_T$
# potential links	$d = 4$	Intensity of choice	$\mu = 8$
Periods	$T = 18$	Independence	$\lambda = 1$
Number of options	$m = 6$	Recency	$\phi_l = \phi_f = 0.9$
Conformity reward	$a_J = 0.05$	Lock-in	$\kappa_l = \kappa_f = 0$
Timing reward	$a_T = 0.15$	Reinforcement	$\delta_l = \delta_f = 1$
Bias on leading	$A_0(0) = 0.6$	Counterfactual	π_0^c by experience

for each agent satisfying the condition of a strongly connected graph.


Two types of figures are employed to display the findings starting with Figure 7 below. For each treatment, the figure is a single representative realization. The first plots the time-series of $\theta_{i,t}$ for the 18 members of the population. The second figure plots the tree representation of the social structure as it is realized in the terminal period of the simulation.

With $(m - 1)a_J \geq a_T$, the default simulation has as its coordinated structure a single leader and the entire population choosing to imitate the leader, directly or indirectly through a chain of imitation links. The variation in the parameters across different treatment only affects the evolutionary process of the population and not the cooperative target.


5.1 Parameter Effects

While multiple players may lead, to be a leader one must have followers. A dominant leader refers to a single leader attracting an outsized share of the population of followers. The simulations presented in this section point to the regular success of the EWA and nested logit to generate a dominant leader in a variety of settings and with a wide range of parameter values.

With the EWA governing individual strategy adjustments over time, an individual's emergence is path dependent. In the early periods of the simulation, success by the individual is the result of random transitory events. In the absence of any strategy adjustment, an outcome favorable to individual i in one period will not likely be repeated in the next. For a leader to emerge, others in the population must respond to an individual's success by decreasing their own $\theta_{i,t}$ and increasing the weight $w_{i,t}^j$. The process of observation and adjustment allows the initially lucky individual to become a successful leader no longer reliant on luck but empowered by her followers. The success of the EWA in generating a dominant leader is in its backwards-looking measure of performance and the induced adjustments in strategy that reward success with more success.



(a) Time-series of $\theta_{i,t}$



(b) Social structure in the final round. Agent #7 emerged as unique leader.

Figure 7: Base (81.3%): $\mu = 8$, $\lambda_f = 1$, $\phi_f = \phi_l = 0.9$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$. A single agent leads in this realization.

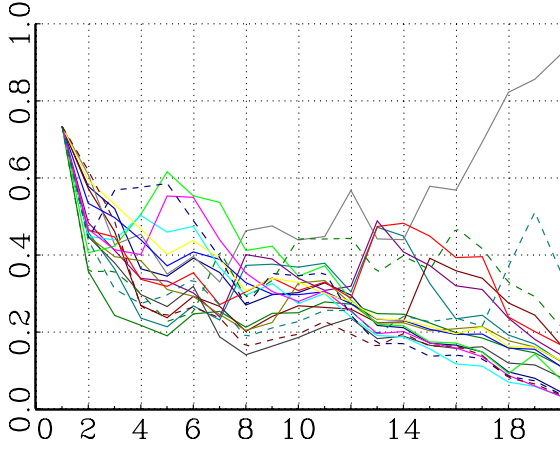
5.1.1 Baseline emergent structure

Figure 7 captures the emergence of a dominant leader and hierarchy of followers from an initially unstructured social setting. As can be seen in 7a, all players start with the same probability of leading. This probability drops off in the early periods as the earnings to the strategy fail to support the initial performance bias assigned to leading. In this example, agent #7 emerges from the population as the only leader with a non-trivial population of followers. Her dominants emerges in round 7. Over time, agent #7's success breeds success becomes embedded in the strategies of the other members of the population. Her emerges as leader benefits the entire population. Of the 1000 iterations, a dominant leader emerges in 813 of the realizations. This frequency is included in the caption of Figure 7.


5.1.2 Intensity of Choice

Manipulating the IOC parameter produces a variety of outcomes. Decreasing the IOC parameter decreases agent sensitivity to performance differences. At a sufficiently low value, individual adjustment to transitory random events are insufficient to create enough of a social advantage upon which to build social structure. The result is a population that remains no more organized than in the initial period.

Figure 8 is produced by a simulation with $\mu = 4$. This is still sufficiently high to enable successful emergence of a unique leader within the time-frame of 20 rounds but the process takes longer and is less frequent. At lower values of μ the process fails completely even when allowed to continue indefinitely. The



(a) Time-series of $\theta_{i,t}$



(b) Social structure in the final period

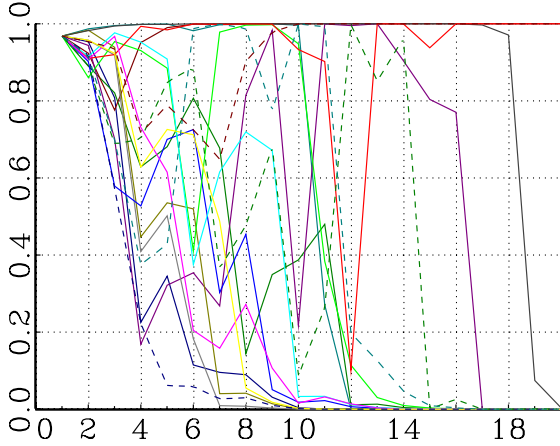
Figure 8: \downarrow IOC (58.5%): $\mu = 4$, $\lambda_f = 1$, $\phi_f = \phi_l = 0.9$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$. A single leader of a unique nontrivial tree emerges. Some members continue to lead. The low IOC slows the process of identifying a leader and collecting followers.

IOC parameter must be sufficiently large to induce an adequate behavioral response to what are initially small earnings advantages that occur when the population is unorganized.


5.1.3 Lock-in

For $\kappa_h = 0$, the performance measures are the average of weighted past performances. For $\kappa_h = 1$, the performance measures are a cumulative measure of past performance. The distinction is that for $\kappa_h = 0$, the expected performance differential is the average of the past performance differences with the weighted sum of the past performances divided by the sum of the weights. For $\kappa_h = 1$, the expected performance differential is simply the sum of the past differences. As an input to the nested logit, the increasing differential drives the probability of adoption to the extremes with the superior action attracting all of the probability weight from inferior actions (if $\phi = 1$). With $\phi = 0.9$, the accumulative impact of $\kappa = 1$ is curtailed by the decay in memory.

Changing κ alone produces changes in the evolutionary process (at least over the 20 period horizon). Relative to the baseline treatment, there is a greater number of members quickly moving to the extremes of always lead or follow in the early rounds so that there is a greater dispersion in the types of behavior. Despite the apparent persistence of a subset of the population towards early leading, eventually a single dominant leader emerges. There is a decrease in the frequency of emergence in the cooperative structure down the 75.7%.



(a) Time-series of $\theta_{i,t}$



(b) Social structure in the final period

Figure 9: \uparrow lock-on (75.7%): $\mu = 8$, $\lambda_f = 1$, $\phi_f = \phi_l = 0.9$, $\kappa_f = \kappa_l = 1$, $\delta_f = \delta_l = 1$

5.1.4 Recency

Decreasing ϕ_l and ϕ_f increases the importance of recent observations on the current respective performance measures. Lowering ϕ increases responsiveness to recent events. This improves the frequency of success in emergence of the coordinated structure to near certainty. It also decreases the time necessary for the structure to emerge, as captured in figure 10. With long memory, once a player emerges as the dominant leader with a population of followers, the structure is stable. The short memory also introduces the possibility of a change in leadership (though not observed in this example).

5.1.5 Independence

Increased independence produced by decreasing λ separates the decision about whether to follow from the decision about whom to follow. As a consequence, while a particular contact may come to dominate the decision about whom to follow, the option to lead may still retain a majority of the probability weight despite being mildly inferior to the dominant contact. This latter feature is reflected in the slower decay in $\theta_{i,t}$ for the non-leaders in the population as seen in Figure 11. The increase in independence also reduces the frequency of emergence of the coordinated structure to about 50%.

5.1.6 Counterfactuals

A decrease in δ reduces the weight that the subjects give to the potential earnings of those actions not employed. The value of δ can be interpreted as capturing the extent to which players observe the untried





Figure 10: \uparrow recency lead (98.1%): $\mu = 8, \lambda_f = 1, \phi_f = \phi_l = 0.5, \kappa_f = \kappa_l = 0, \delta_f = \delta_l = 1$





Figure 11: \downarrow Independence (49.6%): $\mu = 8, \lambda_f = 0.25, \phi_f = \phi_l = 0.9, \kappa_f = \kappa_l = 0, \delta_f = \delta_l = 1$

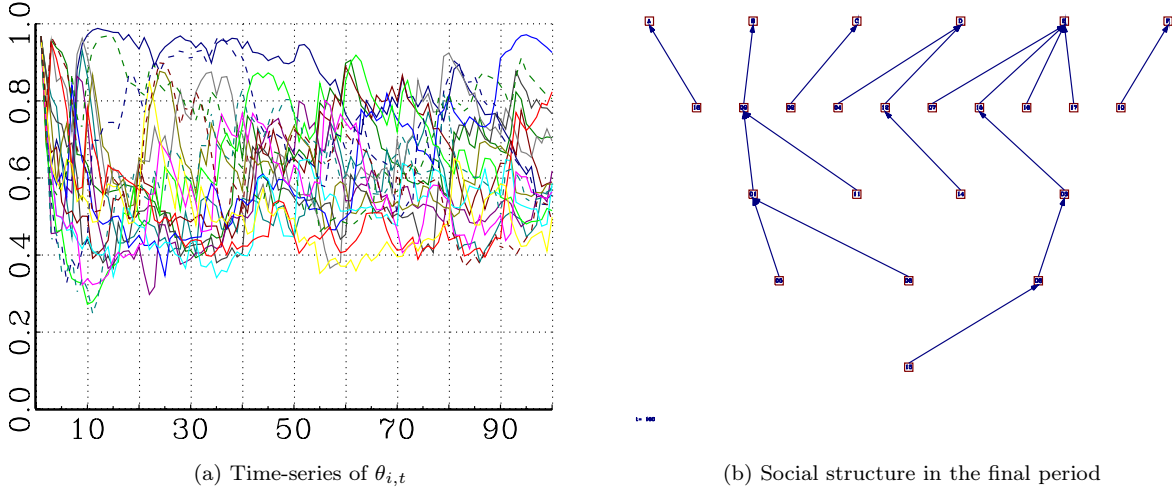


Figure 12: \downarrow Counterfactuals (0%): $\mu = 8$, $\lambda_f = 1$, $\phi_f = \phi_l = 0.9$, $\kappa_f = \kappa_l = 0$, $\delta_f = 0.4$, $\delta_l = 1$

actions. Reducing δ_l means that when a subject follows, they give little weight to the payoff that could have been obtained when leading. The value of $A_{0,t}$ is updated with only a fraction of the potential earning from leading. For $\delta_f < 1$, the player down weights the potential earnings that could have been obtained from the contacts not followed. The down-weighting by player i of contact j occurs whenever player j is not followed regardless of whether player i leads or follows a different contact.

For $\delta_l < 0$ (not shown), the down weighting of the lead counterfactual helps to facilitate the emergence of the coordinated structure. Once a player attempts following, they are effectively less observant of the option to lead and more likely to search out the best following option available.

For $\delta_f < 0$, the players who lead are less observant of the potentially superior earning available when following and when following are less observant of the potentially superior earnings available by following one of the other available contacts. Both of these effects seriously undermine the process that generates convergence on the cooperative structure. Players fail to take advantage of the option to follow and fail to take advantage of the contact offering the highest historic earnings when they do. The consequence can be seen in Figure 12. Of the 1000 iterations of this treatment none succeeded to generate the cooperative structure with $\delta_f = 0.4$.

5.2 Social Advantage

When initial conditions are randomly allocated to the agents, individual characteristics can influence emergence. The ability to be seen by others in the population and potentially imitated is clearly an advantage. A

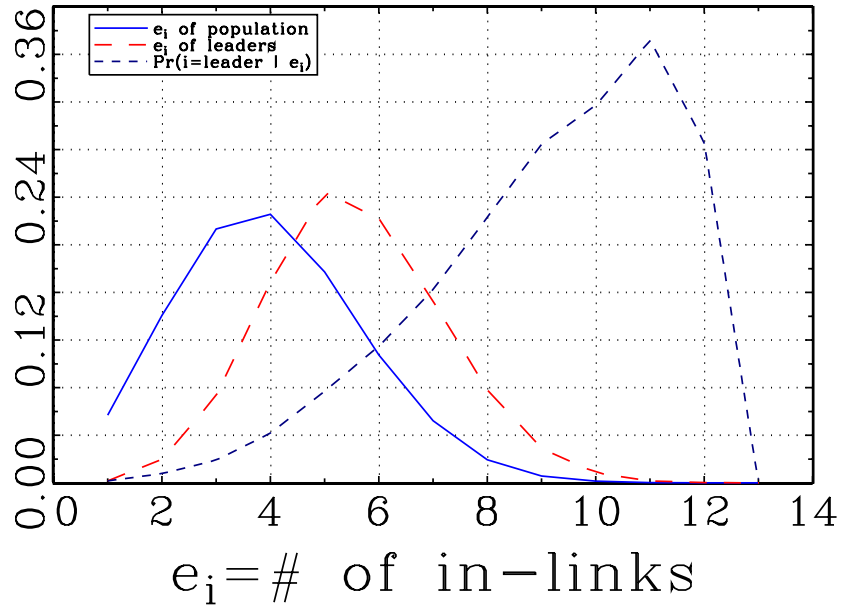


Figure 13: Distribution of e_i among the general agent population (solid blue) and among dominant leaders (long-dashed red). The conditional probability of leading (short-dashed black) increases in e_i .

randomly generated network of contacts produces a non-degenerative distribution of incoming links for each agent. Let e_i be the number of potential links directed at agent i . As revealed in Figure 13, those with a greater number of incoming links have a heightened probability for emerging as a leader. The distributions and probabilities depicted in Figure 13 are generated from 10,000 simulation iterations. The solid (blue) curve is the distribution over the entire population of 180,000 agents as initiated at the start of the simulation. The long-dashed (red) curve is the distribution of the population of dominant leaders as realized at the end of each iteration. The short-dashed (black) curve is the conditional probability of emerging as a leader based on the agent's e_i , peaking at 36% for a player with 11 incoming links, in contrast to the unconditional probability of $1/18$.

The simulations used to generate the distributions in Figure 13 employ the baseline parameters. The rightward shift in the number of incoming directed links for leaders relative to the general population captures the benefit to greater social advantage. A player gains greater observability by being the target of greater number of directed links. This improves the likelihood that favorable transient events will be noticed and accommodated by other players.

Model	Demo(i)	$\hat{\pi}_l^{cf}$	# Contacts			Premium to Lead		Combinations			
			All	3	4	high	low	3&high	3&low	4&high	4&low
Experiment			0.417	0.208	0.556	0.208	0.556	0.167	0.250	0.250	0.708
Global	1	own	0.011	0.011	0.011	0.027	0.0006	0.022	0	0.0325	0.0008
Global	0	own	0.014	0.012	0.015	0.034	0.0003	0.024	0	0.044	0.0004
Global grow	0	own	0.058	0.039	0.071	0.125	0.014	0.077	0.0008	0.172	0.021
Individual	1	own	0.041	0.024	0.052	0.055	0.031	0.045	0.002	0.066	0.045
Individual	0	own	0.074	0.044	0.094	0.077	0.072	0.085	0.003	0.069	0.106
Individual	0	rational	0.088	0.043	0.118	0.066	0.103	0.080	0.007	0.052	0.150

Table 6: Frequency of coordination in virtual experiments. Demo(i) indicate whether the model is estimated using demographic and experience variables. If not, then the estimate of A_{00} is a free parameter. The column $\hat{\pi}_l^{cf}$ identified the model for computing the counterfactual value of leading.

5.3 With condition $(m - 1)a_J < a_T$

A decrease in $(m - 1)\frac{a_J}{a_T}$ changes the incentives to the most distant followers of the dominant leader. Decreasing m improves the likelihood that leading will produce a match with the dominant leader, increasing the expected payoff to the strategy. A decrease in $\frac{a_J}{a_T}$ increases the relative premium reward to leading. Both changes increase the incentive to lead, and for $(m - 1)a_J < a_T$ the incentive is enough to induce the most distant followers of the dominant leader into leading. From (2), the equilibrium size of the follower population is $\mu^* = 6$ as computed based on $a_J/a_T = 6$ and $m = 6$.

6 Virtual Experiments

The experiments suggest that there is a role for both environment and individual decision making in the emergence of the cooperative structure. Simulations of the experiment are used to explore the role of each. The simulations recreate virtually the treatments examined in each of the ten experiment session. The simulations thus have 36 subjects allocated into two groups of 18 participants who play together for 20 periods before being reallocated into a new group of 18. The three different grouping means that the subject plays for 60 periods total. The directed links connecting the subjects are the same as in the actual experiment across all three allocations, as are the rewards for conformity and for timing advantage. The virtual experiments are populated with virtual subjects who adapt their behavior according the the estimated EWA and nested logit parameters.

Table 6 reports the frequency with which the cooperative structure emerges from simulations. The first row reports the frequencies of the actual experiments (as reported in Table 2). Each subsequent row reports the frequency of success from the virtual experiments. The setting of the virtual experiments recreates

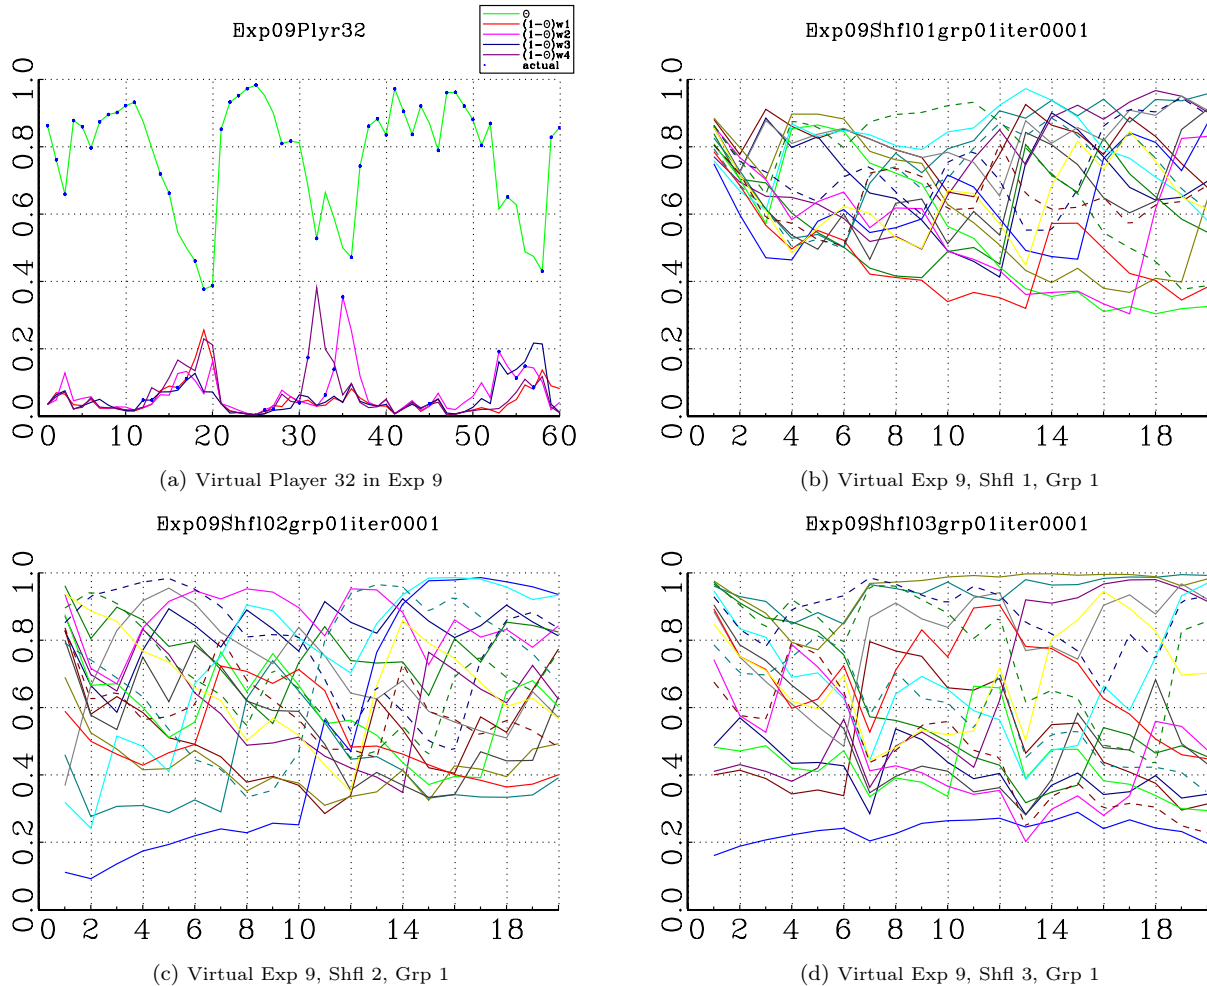


Figure 14: Virtual Experiments: evolution of probability weights with EWA and nested logit parameters from the Global model with individual characteristic information.

the settings of the actual experiments. There are ten different sessions replicating the treatments of the experiments. Each session is iterated 200 times. The different virtual treatments consider different models of individual adaptive behavior using the global and individual model estimates. A feature that improves the ability of the virtual subject to coordinate depends on how they compute the counterfactual value of leading. For $\hat{\pi}_i^{cf}$ identified as “own,” the subjects base performance of leading on the payoff to their own considered choice. For $\hat{\pi}_i^{cf}$ identified as “rational,” the subjects compute the value of leading based on the calculating the average payoff to those who lead.

The first frame of Figure 14 plots the probability weights associated with each action for a player during the same experiment. The virtual subjects are modeled using the global parameters from the model with individual characteristics. The remaining three frames displays the evolution of θ_i for the population in

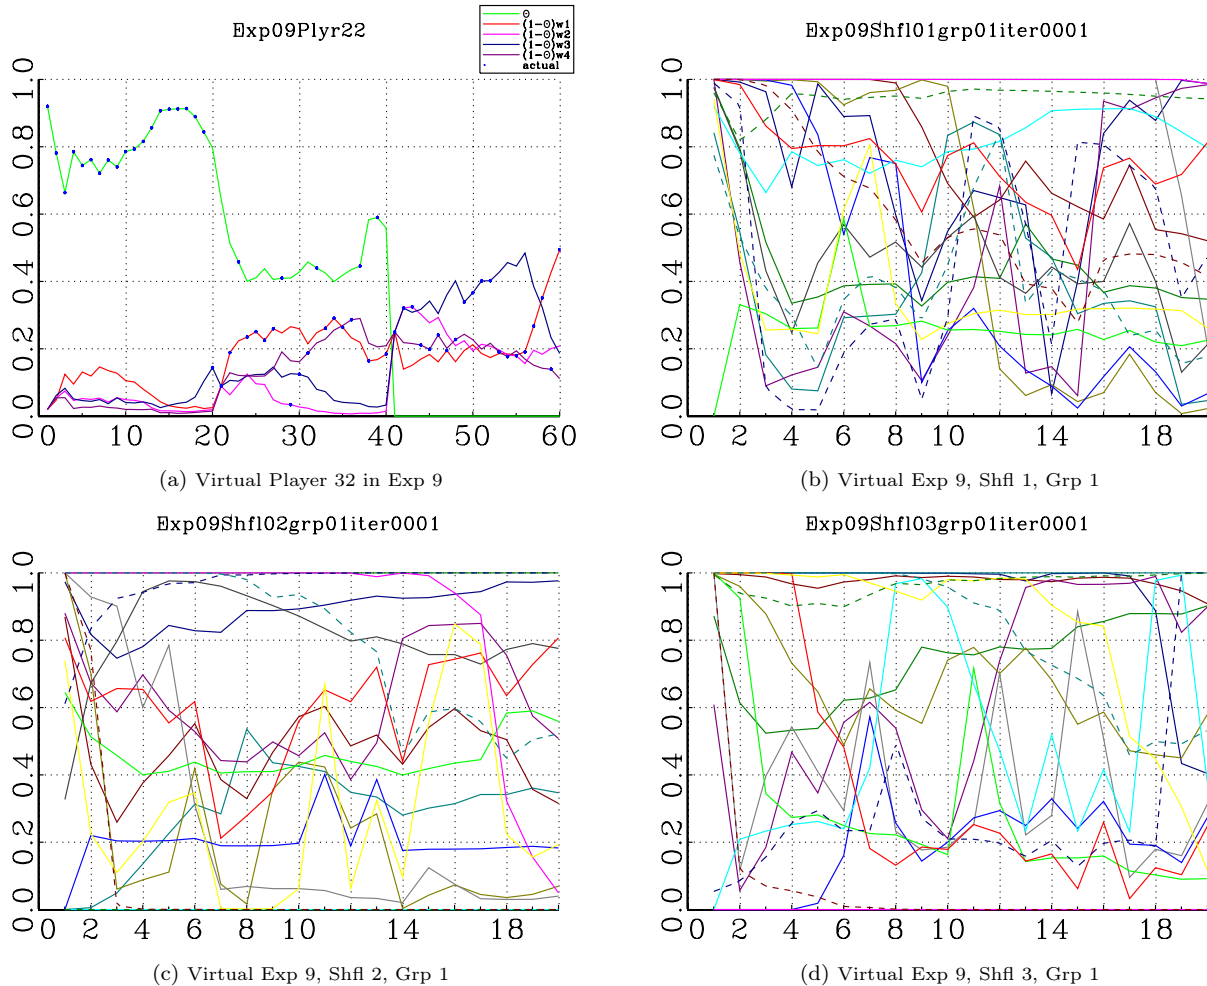


Figure 15: Virtual Experiments: evolution of probability weights with EWA and nested logit parameters from the Individual model with no individual characteristics.

group 1 of each of the three shuffles of iteration 1 of the virtual experiment 9. The evolution in θ_{32} and $(1 - \theta_{32})w_{d,32}$ reflect the failure of the coordinated structure to emerge from the population. There are a number of features of the global parameter estimates contributing to the failure to produce the coordinated structure. The value of μ is too low to generate an adequate response to the superior performance observed superior performance in the follow strategy. In addition, δ_f is also too low, preventing the subjects from recognizing the superior performance offered by the follow strategy when they lead.

For Figure 15, virtual subjects are modeled using the individual parameters from the model without individual demographic information. While still failing to match the rate of success observed in the experiments, these simulations do display greater individual evolution in response to an evolving environment. A dominant leader typically emerges who has a substantial following. Overall, an insufficient number of players

adapts to the emergence of a leader within the given 20 periods of play. There are too many of the virtual subjects who stubbornly continue to lead, undermining the emergence of a leader for the remainder of the population to identify and imitate.

7 Conclusion

With an incentive to cooperate a large population of players is capable of developing the relationships necessary to generate the coordinated outcome. That is, the subjects come to rely on social connections to identify and achieve coordination on each new alternative. Success is facilitated by appropriate rewards that emphasize cooperation over competition and by increased social contacts that increase visibility and decrease the distance that information must travel.

The combined Experience Weighted Attractor and Nested Logit model produces insight into how individuals update strategies based on personal experience and recall. There is considerable heterogeneity among players. Those who successfully emerge as leaders reflect the heterogeneity of the larger population, suggesting that environment and luck play a substantial role in determining who becomes a leader.

Simulations reveal some of the important features that are conducive to the emergence of the coordinated structure. A homogeneous population distributed over a regular network of directed links will converge to the coordinated structure for a broad range of parameters for the EWA and nested logit models. Convergence on the coordinated structure arises from appropriately adaptive behavior by individual subjects responding to the events they can observe. The key is that there has to be sufficient response by the subject to give preference towards actions that reinforce lucky outcomes. Such adjustments will give social advantage based on otherwise transient events. Continued adjustment rewards advantage, transforming initial luck into a permanent advantage, eventually producing a leader and followers. For this to occur requires sufficiently observant subjects, meaning that subjects have to pay attention to the counterfactuals offered by following in order to recognize the benefits of following the right contact. Subject also have to be responsive to differences in performance as they emerge. A population that is slow to respond will not adapt the network to take advantage of transient events before they have passed.

The representative agent estimated from the pooled data does not have the characteristics necessary to produce convergence to the cooperative structure. The value of μ is too low for the agents to respond to the performance differentials that emerge when there is little social organization. Additionally, δ_f is also too low so that the agents are not able to identify the right contact to follow and the advantage offered by

following. The failure of the representative agent model also points to a failure to capture the heterogeneity of the actual subject population.

Modeling individual behavior generates greater frequency in the emergence of the coordinated structure than the representative agent. Greater adaptation by even a small portion of the population provides the seeds of organization that others can respond to. There is clearly room for improving the model of subject behavior within the experiments. The model fails to fully capture aspects of individual behavior that allows human subjects to achieve the coordinated structure with greater frequency than the virtual subjects. This will be the subject of continued effort.

References

- Acemoglu, D., K. Bimpikis, and A. Ozdaglar (2010, September). Dynamics of information exchange in endogenous social networks. Working Paper 16410, National Bureau of Economic Research.
- Amir, R., F. Garcia, and M. Knauff (2010). Symmetry-breaking in two-player games via strategic substitutes and diagonal nonconcavity: A synthesis. *Journal of Economic Theory* 145(5), 1968 – 1986.
- Amir, R. and J. Wooders (1998). Cooperation vs. competition in R&D: The role of stability of equilibrium. *Journal of Economics* 67(1), 63–73.
- Arndt, J. (1967). Role of product-related conversations in diffusion of a new product. *Journal of Marketing Research* 4(3), 291–295.
- Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Review* 84(2), 406–411.
- Bala, V. and S. Goyal (2000). A noncooperative model of network formation. *Econometrica* 68(5), 1181–1229.
- Banerjee, A. V. (1992). A simple-model of herd behavior. *Quarterly Journal of Economics* 107(3), 797–817.
- Brock, W. A. and S. N. Durlauf (2001). Discrete choice with social interactions. *Review of Economic Studies* 68(2), 235–260.
- Camerer, C. and T. H. Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica* 67(4), 827–874.

- Cowan, R. and N. Jonard (2004). Network structure and the diffusion of knowledge. *Journal of Economic Dynamics & Control* 28(8), 1557–1575.
- Crawford, V. P. and H. Haller (1990). Learning how to cooperate: Optimal play in repeated coordination games. *Econometrica* 58(3), pp. 571–595.
- Dixit, A. (2003). Clubs with entrapment. *The American Economic Review* 93(5), 1824–1829.
- Dutta, B. and M. O. Jackson (2000). The stability and efficiency of directed communication networks. *Review of Economic Design* 5(3), 251 – 272.
- Dwyer, P. (2007). Measuring the value of electronic word of mouth and its impact in consumer communities. *Journal of Interactive Marketing (John Wiley & Sons)* 21(2), 63 – 79.
- Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica* 61(5), 1047–1071.
- Galeotti, A. and S. Goyal (2010). The law of the few. *The American Economic Review*, 1468–1492.
- Goldbaum, D. (2014). Equilibrium analysis of conformity and influence on a social network. SSRN Working Paper 1566611.
- Greiner, B. (2004). An online recruitment system for economic experiments. In K. Kremer and V. Macho (Eds.), *Forschung und wissenschaftliches Rechnen*, pp. 79 – 93. GWDG Bericht 63, Ges. für Wiss. Datenverarbeitung, Göttingen.
- Hamilton, J. H. and S. M. Slutsky (1990). Endogenous timing in duopoly games: Stackelberg or Cournot equilibria. *Games and Economic Behavior* 2(1), 29 – 46.
- Heal, G. and H. Kunreuther (2010). Social reinforcement: Cascades, entrapment, and tipping. *American Economic Journal: Microeconomics* 2(1), 86–99.
- Hill, S., F. Provost, and C. Volinsky (2006). Network-based marketing: Identifying likely adopters via consumer networks. *Statistical Science* 21(2), 256–276.
- Jackson, M. O. and A. Watts (2002). The evolution of social and economic networks. *Journal of Economic Theory* 106(2), 265–295.
- Jackson, M. O. and A. Wolinsky (1996). A strategic model of social and economic networks. *Journal of Economic Theory* 71(1), 44–74.

- Katz, E. and P. Lazarsfeld (1955). *Personal Influence: The Part Played by People in the Flow of Mass Communications*. Foundations of Communications Research. Glencoe, Ill: Free Press.
- Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *American Economic Review* 75(3), 424–440.
- Kirman, A., S. Markose, S. Giansante, and P. Pin (2007). Marginal contribution, reciprocity and equity in segregated groups: Bounded rationality and self-organization in social networks. *Journal of Economic Dynamics & Control* 31(6), 2085–2107.
- Reinganum, J. F. (1985). A two-stage model of research and development with endogenous second-mover advantages. *International Journal of Industrial Organization* 3(3), 275 – 292.
- Sadanand, V. (1989). Endogenous diffusion of technology. *International Journal of Industrial Organization* 7(4), 471 – 487.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology* 1(2), 143–186.
- Schelling, T. C. (1973). Hockey helmets, concealed weapons, and daylight savings - study of binary choices with externalities. *Journal of Conflict Resolution* 17(3), 381–428.
- Tapiero, C. S. (2014). A financial ccapm and economic inequalities.
- Tesoriere, A. (2008). Endogenous R&D symmetry in linear duopoly with one-way spillovers. *Journal of Economic Behavior & Organization* 66(2), 213 – 225.
- Watts, A. (2001). A dynamic model of network formation. *Games and Economic Behavior* 34(2), 331–341.
- Whyte, W. J. (1954). The web of word of mouth. *Fortune* 50, 140–143.