

ECONOMICS DISCIPLINE GROUP

UTS BUSINESS SCHOOL

WORKING PAPER NO. 15

November 2013

Follow the Leader: Simulations on a Dynamic Social Network

David Goldbaum

ISSN: 2200-6788

<http://www.business.uts.edu.au/economics/>

Follow the Leader: Simulations on a Dynamic Social Network

David Goldbaum*

October 31, 2013

Abstract

This paper introduces a process of individual adjustment based on private local experiences and observation that allows for the emergence of a global social structure that is the equilibrium to the static follow-the-leader game of Goldbaum (2013). The setting rewards agents for being early adopters of popular products or trends. From simple, myopic, self-serving adjustment based on historic evidence by individuals emerges the the equilibrium social structure consisting of a single choice leader and a population of followers, which, in the static setting would require an unlikely degree of coordination to produce. Individual actions take place in a social context with individuals linked via one-way paths of observation. The strategy by which an agent chooses among the available options evolves over time. Different adjustment emergent processes contribute towards the understanding of the unfolding of events that generate the equilibrium structure.

Keywords: Leader, Dynamic Network, Social Interaction, Consumer Choice, Simulation
(JEL Codes: D85, D71, C71)

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

1 Introduction

According to the American Heritage Dictionary of Cultural Literacy, “a ‘guru’ is a teacher who attracts disciples or followers.”² Guru is a term often employed to connote a person whose opinion is frequently sought and seen as influential on decisions made by a population of followers. Gurus exist in religion, politics, and commerce. When social influences affect the desirability of a choice, then a guru can play a coordinating role that improves social outcomes. The guru position may be desirable for the social rewards associated with the ability to influence decisions and for offering a financial reward if the ability to influence outcomes can be leveraged for rent. In such a case, the guru position enabling the desirable social coordination is also the object of competitive desires of each individual in the population.

This paper examines the dynamic process by which an individual might emerge as an opinion leader from a population of like individuals seeking the same leadership position. Goldbaum (2013) demonstrates that an environment rewarding early adoption of a subsequently popular choice has in equilibrium a social structure consisting of a single leader who coordinates the decisions of a population of followers. The environment, rather than personal charisma or other leadership characteristics, ensures the leader with a population of followers. This work contributes to Goldbaum (2013) by exploring the process by which the equilibrium hierarchical structure can emerge from an initially unstructured population. This dynamic version of the model replaces Goldbaum (2013)’s fully rational, fully informed agents with simple self-serving agents acting on the limited information of their own personal experiences. The process is driven by repeat interactions, learning, and adjustment. Different processes of learning and adjustment are considered. The variety of outcomes produced offer insight into the process of emergence of a leader.

Individual strategies evolve according to the Experience Weighted Attractor (EWA) of Camerer and Ho (1999). The allocation of probability across the possible individual actions is according to a nested logit model of Hausman and McFadden (1984). The experience weighted attractor allows agents to update their strategies according to their individual private experiences and observations. The nested logit allows for the clustering of possible actions consistent with the variance-covariance structure produced by a sequential

¹Columbia World of Quotations. Columbia University Press, 1996. <http://quotes.dictionary.com/Onedoesnotbecomeagurubyaccident> (accessed: March 17, 2011).

²The American Heritage New Dictionary of Cultural Literacy, Third Edition.

decision process when deciding on an action. Both models have strong foundations in evolutionary and discrete choice modeling. The resulting evolution of strategies is found to be broadly capable of generating the equilibrium social structure. This paper develops the dynamic environment investigated through simulation analysis and explores the variety of processes and outcomes as determined by different combinations of the EWA and nested logit model parameters.

Social networks have received considerable attention in economics for explaining aggregate social behavior. In these models, individuals may be influenced globally by the aggregate behavior of the population or locally by their neighbors, a subset of the population who serve as a reference group. This project incorporates aspects of both. The network is a social structure influencing how discrete decisions are made or new products or technologies are adopted in Schelling (1971), Schelling (1973), and Katz and Shapiro (1985). The utility of a choice depends on the choices of those in an agent's peer group. Because of social interactions, products that garner no particular consumer preference and may even be perceived to be inferior can grow to dominate consumer choice.

A number of environments reward individuals for being early adopters of a trend that subsequently becomes popular. Financial or social rewards come to those who can gain a reputation for being a predictor (or setter) of trends. "Wine geeks just love bragging rights. They get kudos from their peers when they get a high-score wine first or get it cheaper" (Bialek quoted in Los Angeles Magazine, 1998(Dec)). Because the agents seek reward through early adoption, timing of an agent's decision must be modeled as an endogenous component of the social network, a clear deviation from the traditional mean field examinations of local interaction. Further, agents are rewarded based on aggregate behavior rather than on the local behavior of the individual's peer group. Brock and Durlauf (2001) also model the individual as seeking to conform to the social norm, but in an environment in which the payoff is based purely on the aggregate behavior.

The traditional examination of social interactions in economic decisions is based on a fixed social network. Regular structures are convenient for deriving aggregate properties of the network. The level of connectedness determines the rate at which information disperses through the network, which can impact the leader and follower structure.

Dynamic models such as Watts (2001), Bala and Goyal (2000), Jackson and Watts (2002), and Kirman et al. (2007a) are concerned with evolution and convergence. Bala and Goyal (2000) examine the evolution of a network in which agents create links with individuals with whom the benefit exceeds the cost of maintaining the link. Each agent i offers a direct link benefit (that is uniform), as well as the benefit gained by providing indirect access to those with whom agent i is linked. Thus, the benefit of linking to an individual is

endogenous to how connected that individual is to others in the population. One stable attracting network configuration is a star formation where one agent serves as a hub through which all agents connect.

Networks are also a vehicle for information transmission. Directed networks (where the link between two agents is established by one of the linked pair) have been used to examine advertising and marketing, as in Dutta and Jackson (2000). Undirected networks (where the link requires the support of both agents in a linked pair) also serve as a mechanism of information dispersion, as in Ellison and Fudenberg (1995) where word-of-mouth communication can lead to conformity in behavior. Informational objective also play a role in network formation in Acemoglu et al. (2010) where agents collect information through social contact to inform an irreversible decision made urgent by discounting. Acemoglu et al. (2013) examine the role of information projectors in shaping public opinion through social connections.

Information, evolution, convergence, and stability are all features of the proposed project but the payoff structure and its dependence on the timing of a decision alter the role played by the social connections, thus altering the process of formation and self organization. The directional path through which information is transmitted gives the network a hierarchical structure with a leader and follower. That the directed links are established by agents seeking information rather than by advertisers looking to push information is another important distinction affecting the formation of the network and its employment by the agents. The developed model shares a number of feature with Koenig (2012) where agents choose to form links with an exogenously chosen subset of a population to gather information. Noise in the information transmission plays a role in the network formation.

Like this examination, Chang and Harrington (2007) consider an evolutionary network driven by an experience weighted attractor. Their population is heterogeneous with some better suited to innovation and other better suited to dissemination of ideas. They examine the endogenous social network that arises to connect these different types of researchers. The resulting network produces a symbiotic relationship between innovators and those who are able to facilitate in the propagation of the innovator's ideas. For the current investigation, agents must be allowed to adapt to their changing local environment in a payoff improving manner. Simulations are employed to characterize the behavior of the population.

2 Model

Consider a setting in which agents are rewarded for being early adopters of subsequently popular trends. Faced with a set of options from which they can choose only one, agents can make their selection immediately

or they can wait for indicators regarding others are doing before making their decision. Goldbaum (2013) develops such an environment for a static game. The players implement a strategy of immediately choosing one of the options directly or specify a member of the population to imitate. For completeness, I introduce the environment, the possible actions of the agents, and the payoff and refer the reader to Goldbaum (2013) for the formal development of the environment and analysis of the game theoretic equilibrium solutions. The current paper introduces repeated play of the game by a given population and a dynamic process by which the agents can update their strategy.

Let $N = \{1, \dots, n\}$ be a set of agents. Let $O(t) = \{O_{1t}, O_{2t}, \dots, O_{mt}\}$ be a new set of options available in time period t . In each time period, $t \in \{1, \dots, T\}$, each agent looks to adopt from one 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 $f_{i,t}$, determined by nature, map from labels to options, $f_{i,t} : K \rightarrow O(t)$. Each player thus privately observes the K set of labeled options. Each player sees different labels and for every i, j pair there is a one-to-one function $h_{i,j,t} : K_{i,t} \rightarrow K_{j,t}$ that is unknown to the players. The mechanism prevents coordination between agents on a particular choice, either through prior communication or by identifying a focal point.

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 $a_{i,t} = f_{i,t}^{-1}(O_{k,t})$. 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 *imitate* or *follow* the other player. Thus, the set of actions for player i is $A_i = K_i \cup N \setminus \{i\}$. Write $a = (a_1, \dots, a_n)$, for an action profile, where $a_i \in A_i$.

Example 1. As illustrated in Figure 1, consider a population of $n = 7$ agents facing $m = 3$ options. Suppose $a_t = (“A”, 1, 1, 3, “C”, 7, 6)$. That is, both agents 1 and 5 lead with agent 1 choosing the option labeled “A” in her list of the options while agent 5 chooses the option labeled “C”. Agents 2 and 3 imitate agent 1, agent 4 imitates agent 3, agent 6 imitates agent 7, and agent 7 imitates agent 6. These actions determine the social structure for the period. Only agent 1 is a leader of a non-trivial population. Having imitated a leader, directly or indirectly, agents 2, 3, and 4 are followers. Agents 6 and 7 are within a self-referencing loop. They go through the period without adopting one of the options. Frames 1a and 1b in Figure 1 depict the same social structure with different realizations of the mapping from individual labels to options. In 1a, agents 1 and 5 select different options according to nature’s determination of $O_1 = f_1(A)$ and $O_3 = f_5(C)$. In 1b, agents 1 and 5 select the same options according to nature’s determination of $O_2 = f_1(A) = f_5(C)$.

Allow that i is a predecessor of j if j imitates i , either directly or indirectly through a chain of imitators. Let $d_{i,j,t}$ represent the distance from agent i to agent j measured in the number of links that separate them.

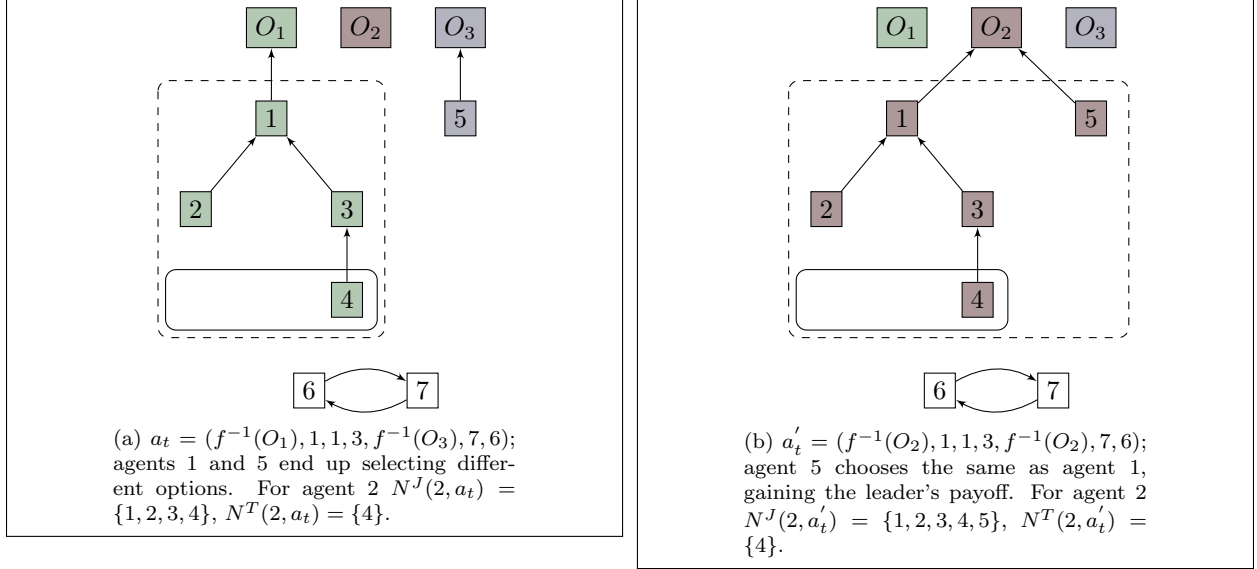


Figure 1: Choice and social organization in a population with $n = 5$, $m = 3$. Payoffs depend on an agent's relative distance from the option selected. An arrow from an individual directly to one of the choices represents a choice by a leader. An arrow from one agent to another represents a path of imitation with the arrow originating from the imitator.

From Example 1, $d_{2,1,t} = d_{3,1,t} = 1$ and $d_{4,1,t} = 2$. An agent's distance from the option she adopts can likewise be measured. Let $d_{i,t}$ represent the distance from agent i to the option she adopts. Starting from a distance of zero for leaders, $d_{1,t} = d_{5,t} = 0$, $d_{2,t} = d_{3,t} = 1$, $d_{4,t} = 2$, and $d_{6,t} = d_{7,t} = +\infty$.

All strategies are decided upon and implemented at the start of the period. Each imitation involves a time delay within the period. As a result, all agents of equal distance from their option adopt their option simultaneously and before those who are of greater distance. In this manner, distance is a substitute measure of the passage of time within a period.

Let $N^J(i; a)$ denote the set of players who adopt the same option as does player i (inclusive of i). Let $N^T(i; a)$ denote the subset of $N^J(i; a)$ who are of greater distance from the option than is i . Let $\mu_{i,t}^J$ and $\mu_{i,t}^T$ represent the size of the respective $N^J(i; a)$ and $N^T(i; a)$ sets. The payoff for player i is

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

with coefficients $a_T \geq 0$ and $a_J \geq 0$. The first element of the payoff is the "conformity" component, much like the community effect of Blume and Durlauf (2001). The second element in (1) reflects the distance/timing advantage a player has over other players. By (1), an individual's payoff is purely a social phenomenon. The options themselves offer no direct benefit to the agent.

The dashed box included in both frames of Figure 1 captures $N^J(2, a_t)$, the agents adopting the same choice as agent 2. The solid box captures $N^T(2, a_t)$. Agent 2's payoff reflects that she adopts the same option in advance of agent 4, even though agent 4 is not a successor of agent 2. In Figure 1a, agent 5 is excluded from $N^J(2; a_t)$ because she adopts a different option. In Figure 1b, she is included based on both players adopting the option O_2 . For agent 5, the realization depicted in 1a yields a payoff of zero while the realization in 1b yields $\pi(5, a'_t) = 5a_J + 3a_T$.

Each agent has a set of d "contacts" who they can unilaterally decide to imitate. Graphically, to imitate is to establish a directed link extending from the agent who decides to imitate to the agent they have decided to imitate. The graph of the "potential" links based on the contact lists is presumed to be strongly connected, meaning that for any two agents, i and j , either there exists a potential link from i to j or there is a chain of potential links that would make it possible for i to imitate j indirectly.

Autarky, the consequence of everyone leading, generates expected payoffs of $a_J(n-1)/m$. At the other extreme, a social order consisting of a single leader who's choice then disseminates through the tree of followers generates a range of payoffs with $(a_T + a_J)(n-1)$ earned by the leader and $a_J(n-1)$ going to the most distant member(s) of the follower tree. The benefit over autarky to even the lowest compensated agent illustrates the need to coordinate. The premium earned by the leader may generate competition between individuals seeking to move from an initially unstructured social network towards a structure offering greater rewards.

As developed in Goldbaum (2013), for

$$B := (m-1) \frac{a_J}{a_T} - \left(1 - \frac{1}{n-1}\right) \geq 0,$$

the structure consisting of one player leading the population of followers is the Nash Equilibrium. The equilibrium retains the leader-follower structure for $B < 0$, but the most distant agents from the leader are better off also choosing one of the options directly. These agents lead, but without any followers. They benefit from the $1/m$ probability of choosing the same option adopted by the leader and earning the leader's payoff, as agent 5 does in Figure 1b. The most distant successors are attracted to lead when the premium to acting in advance of followers is high (so that a_J/a_T is low) or when m is small so that the probability of choosing the same as the leader is high. This situation is reflected in the condition $B < 0$.

The Nash equilibrium to Goldbaum's static game identifies a stable social structure in which no individual can improve their position by unilateral deviation. What the static model fails to offer is a mechanism or

path by which the players can achieve the organized social structure. From a population of n players, only one can be the leader. In addition, each player has to adopt the strategy optimized according to who in the population leads. It is entirely unlikely that the equilibrium social structure emerges in a single shot game. Simulations are employed to investigate dynamic behavior and adjustment to explore possible paths of individual behavior and group organization.

For the dynamic simulation, each agent chooses from the available actions probabilistically. In period t , for $\theta_{i,t} \in [0, 1]$,

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

where $\sum_j w_{i,t}^j = 1$. For agent i , if j is not in her contact list, then $w_{i,t}^j = 0$. Otherwise $w_{i,t}^j \in [0, 1]$. For those agents who lead, a simple random assignment for determining $k_{i,t} \in O_t$ captures the inability to coordinate on an initial choice. For $h \in \{1, \dots, m\}$ and $i \in \{N \setminus s_{i,i} = 1\}$, $\Pr(k_{i,t} = O_h) = 1/m$.

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,

$$\begin{aligned}A_{0,t+1} &= \frac{\phi_l N_{t,l} A_{0,t} + (\delta_l + (1 - \delta_l)\mathbf{1}(x_t = l))\pi_{l,t}}{N_{t+1,l}} \\ N_{t+1,l} &= 1 + \phi_l(1 - \kappa_l)N_{t,l}\end{aligned}\tag{2}$$

and for following agent j ,

$$\begin{aligned}A_{j,t+1} &= \frac{\phi_f N_{t,f} A_{j,t} + (\delta_f + (1 - \delta_f)\mathbf{1}(x_t = d))\pi_{j,t}}{N_{t+1,f}} \\ N_{t+1,f} &= 1 + \phi_f(1 - \kappa_f)N_{t,f}.\end{aligned}\tag{3}$$

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

For performance measures A_0, A_1, \dots, A_d ,

$$\Pr(i \text{ leads}) = \theta_{i,t} = \frac{\exp(\mu A_{0,t})}{\exp(\mu A_{0,t}) + \left(\sum_{j \in D} \exp(\mu A_{j,t}/\lambda)\right)^\lambda} \quad (4)$$

$$\Pr(i \text{ follows } j) = (1 - \theta_{i,t})w_{i,t}^j = \frac{\exp(\mu A_{j,t}/\lambda) \left(\sum_{j \in D} \exp(\mu A_{j,t}/\lambda)\right)^{\lambda-1}}{\exp(\mu A_{0,t}) + \left(\sum_{j \in D} \exp(\mu A_{j,t}/\lambda)\right)^\lambda}. \quad (5)$$

As shaped by the parameters of the EWA, the performance of each strategy is updated based on the most recent realized payoff or counterfactual. 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”.

In empirical studies, a nested logit is required when there is correlation between the choices in a multinomial logit environment. 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. In a simulation setting the nesting serves as a mechanism for compartmentalizing components of a decision.

The parameter μ is commonly referred to as the intensity of choice (IOC) parameter. It determines how sensitive an agent is to differences in the performance measure. At $\mu = 0$, the agent is indifferent to the choices regardless of the performance. As μ increases, for a fixed performance differential, the agent’s probability of adopting the superior strategy increases.

The λ parameter controls the extent to which the agent 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 option 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 other 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.

A variety of alternatives to the combined EWA and nested logit are available for dynamically modeling strategy adjustment. Also modeled, but not included for presentation in this paper, were two versions of a two-stage decision process with separate EWA processes for determining the value of $\theta_{i,t}$ and for determining the value of $w_{i,t}^j$. The two-stage decision was investigated using both the Camerer and Ho (1999) specified power distribution for allocating probabilities based on the performance measures as well as an alternate method based on a k -choice replicator dynamics derived from Branch and McGough (2008). The latter was particularly useful in its ability to consistently produce the same equilibrium-consistent social structure from different randomly generated starting conditions.³ The EWA and nested logit combination presented here is more sensitive to initial conditions so that a given set of parameters produce more varied simulation outcomes. The advantage of the EWA and nested logit combination is that the models are well grounded empirically and has been used extensively for estimating human behavior in a wide range of discrete choice settings, features absent from the two-stage decision models. The two-stage decision process, while intuitively appealing in many applications, typically cannot be employed to empirically estimate a sequential k -choice discrete choice as its structure is typically inconsistent with the underlying data generating process. A straight application of the Camerer and Ho (1999) model as a single-stage decision performed poorly because of the inability to effectively separate the IOC in decision about whether to lead from the decision about who to follow. As a result, it was difficult, with one parameter, to generate reasonable behavior in both the leader with a population of followers and among the follower population.

Once actions have been determined for the period, the action-dependent social network structure can be constructed and payoffs computed according to (1). The EWA updates the performance measures of the untried options as well as the option employed. For the former, the agents need to estimate the payoff that would have been earned with each of the counterfactual actions. *Ex-post*, the agents can observe which option each member of their contact list adopts and when within the period this adoption takes place. With knowledge about how popular each option is over time, the agents can place themselves in the position of having imitated each contact and guess at the payoff that would have been earned as the basis for the counterfactual to the actual action taken. What an agent does not know is which or how many of the subsequent adoptions of their choice can be attributed to their own successors. Were an agent to switch strategies, any successors would consequently switch with them. Rather than attempt to speculate on this

³This made it particularly useful for guiding the development of the theoretical model presented in Goldbaum (2013)

Table 1: Base Parameters

Population	$n = 100$	Participation	$B = 0.11 > 0$
# potential links	$d = 6$	Intensity of choice	$\mu = 1$
Periods	$T = 500$	Independence	$\lambda = 1$
Number of options	$m = 12$	Recency	$\phi_l = \phi_f = 1$
Conformity reward	$a_J = 0.1$	Lock-in	$\kappa_l = \kappa_f = 0$
Timing reward	$a_T = 1$	Reinforcement	$\delta_l = \delta_f = 1$
Prior on leading	$A_0(0) = 0$	Counterfactual	π_0^c by experience

aspect of the social network, the agents simply take the network, except for their own position, as given.

Each agent, whether leading or following, chooses one option as their direct choice for that period should they choose to lead. Those that follow use the payoff that would have been obtained had the agent led as the experience-based estimate of the counterfactual to leading.

3 Simulations

Table 1 reports the baseline parameters of the simulation. A random process determines the contact list for each agent satisfying the condition of a strongly connected graph.

One measure of how close the population has come to adopting equilibrium strategies is to look at the number of agents who lead. Let $\mu_{lead,t}$ be the number of agents who lead in period t . For $B \geq 0$, $\mu_{lead,t} = 1$ in equilibrium. Another measure of success is to consider each agent's distance from their chosen leader. A minimal distance social structure is one in which each follower chooses to rely on the contact offering the shortest distance to her leader. Having already defined $d_{i,j,t}$ as agent i 's distance from the chosen leader in period t , let $d_{i,j,t}^*$ indicate the the shortest possible distance. Let Δ_t , measured aggregate deviations from minimal distance according to

$$\Delta_t = \sum_i d_{i,j,t} - d_{i,j,t}^*. \quad (6)$$

Accordingly, $\Delta_t = 0$, referred to as Δ -efficient, represents a social structure in which each follower employs the shortest distance to her leader. Values of $\Delta_t > 0$ indicates deviations from minimal distance. A third related measure to the proximity to the equilibrium structure is how many in the population are caught in a self-referencing imitation loop. Let $\mu_{loop,t}$ represent the number of agents caught in a self-referencing loop.

Three types of figures are employed to display the findings starting with Figure 2 below. The first plots the time-series of $\mu_{lead,t}$, $\mu_{loop,t}$, and Δ_t . A second time-series figure plots the values of A_0, A_1, \dots, A_d for the indicated agent in the population. The third figure plots the tree representation of the social structure as it

is realized in the terminal period of the simulation. Across the top of the figure are the m choices, labeled using capital letters. Individual agents, labeled using numbers, appear in rows below the choices based on the distance from the adopted option.

With $B > 0$, the default simulation has as its equilibrium 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 leaves B unchanged and thus only affects the evolutionary process of the population and not the equilibrium target.


3.1 Parameter Effects

While multiple leaders may be present within a population, 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. The equilibrium structure consisting of a unique leader when $B \geq 0$ is more difficult to achieve, demanding particular parameters producing the needed environment such that all but the dominant leader consistently follow using the Δ -efficient link to connect to the leader.


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 the lucky individual's success by increasing the weight, $w_{i,t}^j$, if they are linked to someone successful and decreasing their own $\theta_{i,t}$ in response to the higher payoff offered by imitation. 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.

3.1.1 Baseline emergent structure


Figure 2 captures the emergence of a dominant leader and hierarchy of followers from an initially unstructured social setting. In this example, agent #27 emerges from the population as the only leader with a non-trivial population of followers. Over time, agent #27's success breeds success for herself and for the entire population. As reflected in agent #100's increasing performance measures plotted in frame 2b, payoffs increase along with the number of successors to #27. The nearly unanimous adjustment by the remaining



(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.



(b) Time-series of performance measures as experience by Agent #100 who's shortest distance to the leader is two links.



(c) Social structure in the final period ($t = 500$). Near Δ -efficient hierarchy based on Agent #27 as leader.

Figure 2: Base: $\mu = 0.2$, $\lambda = 1$, $\phi_f = \phi_l = 1$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$. A dominant leader emerges but in the presence of a small persistent population of additional leaders.

population towards following agent 27 makes the structure close to that suggested by the static equilibrium. In deviation from the equilibrium structure, a small but persistent number of agents lead. Additionally, the Δ -efficiency score fails to reach zero, indicating that some followers do not occupy their personal shortest distance to agent #27 when following.


Both deficiencies arise from the same root cause: inadequate differentiation between the different performing options that leaves the inferior actions with non-zero probability weight. The problem is most pronounced for those who are most distant from the leader as the payoff differentials across their options are smaller. With non-zero probabilities on inferior actions, they end up choosing to imitate a contact offering a longer chain of links to the leader or choose to lead rather than follow. For the most distant followers of agent #27, the probability of leading stabilizes at about 20% for the current $\mu = 0.2$. As a consequence, there is a changing but persistent population of agents who lead despite their lack of personal successors with leading and the lower expected payoff it offers. Figure 2b plots the values of the performance measures over time for agent #100, whose distance from agent #27 is $d_{100,27} = 2$. The jaggedness observe in A_{lead} is the product of using the experience driven approach to computing the counterfactual value to leading. Each upward spike is an instance in which agent #100's choice, had he led, would have matched the choice of agent #27.

3.1.2 Intensity of Choice

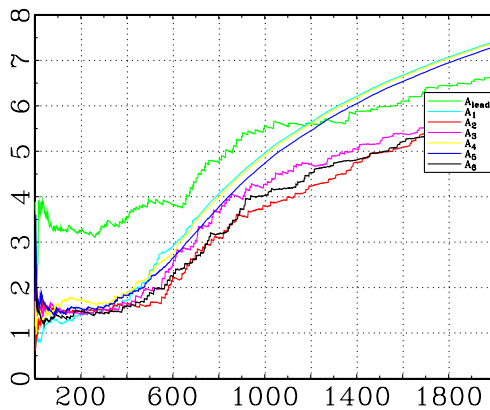
Manipulating the IOC parameter produces a variety of outcomes. Decreasing the IOC parameter below $\mu = 0.2$ decreases agent sensitivity to performance differences. At a sufficiently low value, individual adjustment to transitory random events are insufficient to create enough of a social advantage upon which to build social structure. The result is a population that remains no more organized than in the initial period.

Increasing the IOC to $\mu = 0.5$ increases agent sensitivity to performance differences. Initially, this reduces the number of leaders and brings the Δ -efficient score closer to zero, but fails to achieve $\mu_{\text{lead},t} = 1$ and $\Delta = 0$. As seen in figure 3, the increase in the IOC causes the population to quickly produce near autarky before slowly organizing around one dominant leader. The simulation is run for 2000 periods rather than the default 500 periods. It is not until around period 1200 that agent #52 learns that following agent #100 (at a distance $d_{52,100} = 3$) is superior to leading. Allowing the simulation to run or 10,000 periods still fails to achieve $\mu_{\text{lead},t} > 1$.


Further increasing to the IOC, to around $\mu = 3$, results in the population remaining stuck at near autarky. Adjustment to the IOC parameter alone can generate a social structure very close to the equilibrium structure



(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.



(b) Time-series of performance measures as experience by Agent #52.



(c) Social structure in the final period ($t = 2000$).

Figure 3: \uparrow Intensity of Choice: $\mu = 1$, $\lambda = 1/3$, $\phi_f = \phi_l = 1$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$. Agent #100 emerges as the dominant leader but the high IOC slows the process of collecting followers.

but not the equilibrium.

3.1.3 Lock-in

Increasing lock-in adjusts the perception of performance from one measured in the time-series average to one based on the cumulative history. The former produces performance differentials that stabilize to reflect the average difference in performance between actions while in the latter the performance difference accumulates over time. As a consequence, increasing lock-in by increasing κ amplifies the agents' perception of performance differentials over time.

Let $\pi_{h,t}^\delta$ reflect the performance of strategy h in period t after discounting counterfactual estimates according to δ_h . Thus,

$$\pi_{h,t}^\delta = (\delta_h + (1 - \delta_h)\mathbf{1}(x_t = h))\pi_{h,t}.$$

The impact of lock-in on the performance measure can be seen by examining the formula for the realized $A_h(t+1)$, where, for $h \in \{0, 1, \dots, d\}$,

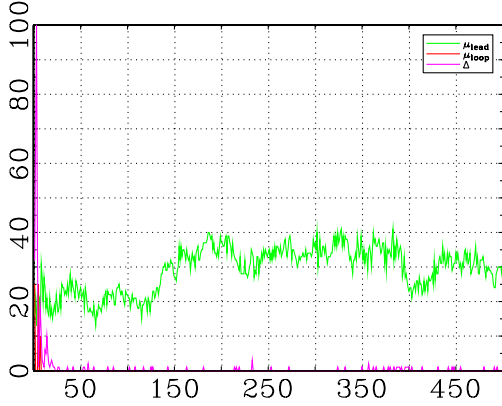
$$A_h(t+1) = \left(\sum_{s=0}^{t-1} (\phi_h(1 - \kappa_h))^s \right)^{-1} \left(\sum_{s=0}^{t-1} (\phi_h^s \pi_{h,t-s}^\delta) \right). \quad (7)$$

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.

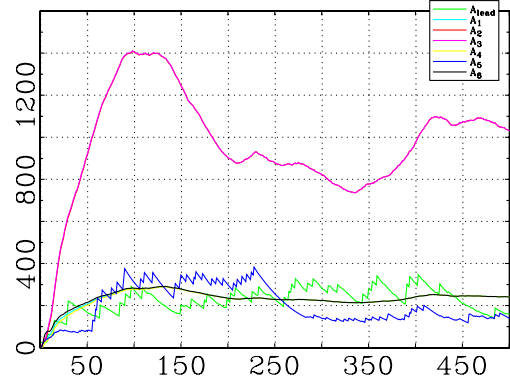
For a $\pi_{h,t}$ drawn from a stationary distribution with mean $\bar{\pi}_h$, then for either $\phi < 1$ or $\kappa < 1$, $A_h(t)$ converges to a stationary distribution around a fixed mean value,

$$\bar{A}_h = \frac{1 - \phi_h(1 - \kappa_h)}{1 - \phi_h} \bar{\pi}_h. \quad (8)$$


The combination of $\phi = 1$ and $\kappa = 1$ ensures that a consistently perceived dominating action eventually attracts all of the probability weight. For $\delta_h = 1$, perception conforms with the actual realized performance. For $\delta_h < 1$, examined in sub-section 3.1.6, perception is biased by reinforced learning. From (8), unequal κ_h



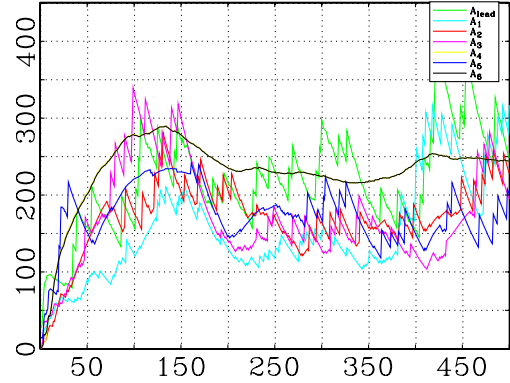
(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.



(b) Time-series of performance measures as experience by Agent #100.



(c) Social structure in the final period ($t = 500$).



(d) Time-series of performance measures as experience by Agent #99.

Figure 4: \uparrow Lock-in with experience-based counterfactual: $\mu = 0.1$, $\lambda = 1$, $\phi_f = \phi_l = 0.98$, $\kappa_f = \kappa_l = 0.8$, $\delta_f = \delta_l = 1$. Agent #10 emerges as dominant but with a large a persistent population of non-followers. Agent #100 imitates #10 from a distance of two links. Other options typically dominate that of following #10 from a distance of three links or greater.

generates an asymptotic bias in favor of the option with $\kappa_h = \max(\kappa_h)$. Only simulations based on $\kappa_l = \kappa_f$ are presented.

The infinitely cumulative differentiation of performance produced by $\phi = 1$ and $\kappa = 1$ cannot be employed in even finite time as the increasing values of A_h quickly produce realizations of machine measured infinite values in (4) and (5). Nonetheless, $\phi \rightarrow 1$ allows for finite accumulated differentiation with the potential of addressing the failure to generate sufficient differentiation to eliminate the use of inferior strategies as seen in the baseline simulation.


Changing κ alone also fails to produce the equilibrium structure. As observed in Figure 4, a substantial portion of would-be followers of the dominant leader (agent #10) either act independently or form small

following trees independent of agent #10. The explanation for this deviant behavior is rooted in the interaction of two features of the environment. First, as reflected in frames 4b and 4d plotting performance measures, a long memory is not the same as an infinite memory so that the steady-state performance values are achieved by period 100. As a result, transitory events can overwhelm long-run features so that an agent may end up chasing after short-run advantages.


The transitory events that sustain the population of non-dominant leaders are the product of the experienced-based computation of the lead counterfactual. Recalling that each follower has a hypothetical selection of one of the options used to compute the counterfactual, luck in the option choice that coincides with the dominant leader’s choice produces a spike in A_0 . The approximately 1 in m followers choosing the same as the dominant leader helps sustain the large population of non-dominant leaders. A secondary effect is that each time one of the non-dominant leaders actually matches choice with the dominant leader, the high payoff received attracts the more distant followers to switch leaders. Over time, the non-equilibrium strategies fail to pay off and the non-dominant agents return to following the leader but the system has stabilized to such that average of new defectors is equal to the average rate at which defectors returning to the equilibrium strategy.

The condition driving the behavior of agents #99 and #100 helps illustrate. From 4b, the superiority to agent #100 of choosing to imitate her personal contact number three is apparent early in the simulation. Contact three offers agent #100 the Δ -efficient distance to agent #10, $d_{100,10}^* = 2$. For agent #99, her contacts four and six both offer the Δ -efficient distance of $d_{99,10}^* = 3$. These options offer consistently high and historically smooth performance as captured in Figure 4d however these Nash equilibrium actions are often dominated by other actions temporarily offering superior performance. Were memory infinite, the consistent high performance would accumulate to dominate the option set. Those choosing directly and lucky enough to choose the same as agent #10 attract followers based on the high they payoff receive. It takes time before an agent following a leader other than agent #10 realizes sufficient loss from this strategy to induce a switch, only to make the same mistake with another attempted leader. These followers of non-dominant leaders perpetuate the leading population by increasing the expected payoff to leading. This increases the number of agents attempting the strategy and lengthening the time that they attempt it.

The experience-based approach may be seen as naive in the current setting. The agent behaves as though the random coincidence of matching the dominant leader indicates an ability to pick future trends. A more sophisticated agent recognizes the random aspect to matching with another leader and appropriately computes the counterfactual value of leading by taking the average payoff produced by leading with each



(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.



(b) Social structure in the final period ($t = 500$).

Figure 5: Lock-in with expectations-based counterfactual: $\mu = 0.1$, $\lambda = 1$, $\phi_f = \phi_l = 0.98$, $\kappa_f = \kappa_l = 0.8$, $\delta_f = \delta_l = 1$. The Nash equilibrium structure quickly emerges

choice. The sophisticated agent’s expectations-based approach smooths the A_0 time-series. As revealed in Figure 5, the absence of the spikes, which draws followers into leading, allows the Nash equilibrium to emerge.

It would appear from this the simulations that sufficient parameters to generate the Nash equilibrium structure identified in Goldbaum (2013) from an initially unorganized population are that (i) $\mu > 0$ so that agents respond to performance differences, (ii) $\kappa > 0$ so that the superior performing individual actions that produce the equilibrium structure can, over time, attract all of the probability weight, and (iii) the environment sufficiently resilient to shocks which might generate a self-sustaining population of agents pursuing inferior strategies. The example, with its quick convergence to the equilibrium structure, is typical but not uniform in repeated simulations employing the same parameters. Of 200 iterations, 162 produced the first realization of a single leader within the first 100 periods while 19 had more than one leader in the terminating 500th period. Within these 19 non-Nash producing simulations, the average number of leaders in the final period was 41 with a minimum observation of 2 and a maximum of 86.

3.1.4 Recency

Decreasing ϕ_l and ϕ_f increases the importance of recent observations on the current respective performance measures. Small declines from one in either ϕ produce no discernible difference relative to the baseline simulation. Lowering ϕ further leaves in place the ability to generate a dominant leader but introduces instability in the social structure, reducing the security of the dominant leader’s hold on the position. Two


distinct narratives of different processes emerge, depending on how the counterfactual value to leading is estimated.

For $\phi_l = \phi_f = 0$, the performance measures reflect only the most recent observation. As a consequence, κ drops out of the EWA. For Figure 6, followers employ the more naive experience-based method when computing the counterfactual for leading. Even in this setting dominant leaders emerge but they rise and fall in quick succession. Agent #77, observed in Figure 6c as the dominant leader, emerged into the position only 15 periods before the end of the simulation. Agent #77 was preceded by a long line of prior dominant leaders. A review of the transition from one dominant leader to the next reveals that a leader typically loses her dominance through the attrition of followers. Again, the experience-driven approach to the lead counterfactual contributes to a large population of leaders. Short memory and random events generate frequent changes in the social structure such that followers abandon a particular leader after only a short period of following. The insufficient time between changes for the followers of the dominant leader to organize contributes to the difficulty in retaining followers. Before the Δ -efficient structure emerges, the more distant followers are attracted to a competing leader offering shorter distances.

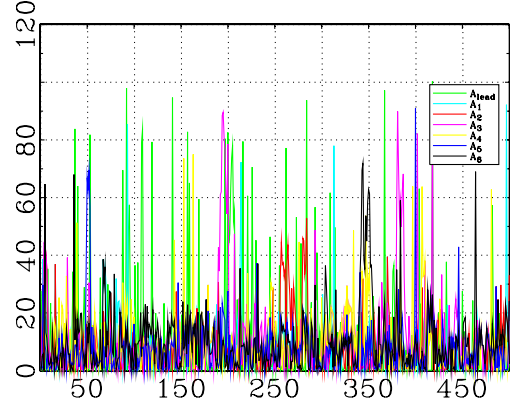
The stability in the A_0 performance measure generated by use of the expectations-based counterfactual helps stabilize the social structure. From Figure 7a, only a small number of leaders attract followers. While the dominant leader in the terminal period is agent #61, agent #74 holds the dominant leader position for the majority of the simulation. She loses the position in period 367 not through attrition of followers but by relinquishing the position to agent #71. Agent #71 is in agent #74’s contact list. In period 366, agent #71 leads and chooses the same option as does agent #74. Agent #74’s estimated counterfactual payoff to following #71 is less than the realized payoff to leading, but not by very much. Short memory in the EWA in combination with small payoff differential leads the nested logit to produce only a small probability advantage to leading over following for both agents #71 and #74. In period 367, #71 leads again and #74 imitates #71. The entire population adopts new actions to optimize based on #71 as the dominant leader.

3.1.5 Independence


Decreasing λ increases the distinction between the decision regarding whether to lead and the decision about who to follow. Recall that for $\lambda < 1$, the leading is evaluated against an aggregate “follow” action rather than against each of the individual $d = 6$ acts of imitating a particular individual the the agent’s contact list. As a consequence, the initial distribution of the population consists of a greater number of leaders than observed in the reference simulation. As the population evolves, the distribution of probability weight towards the



(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.




(b) Time-series of performance measures as experience by Agent #100.

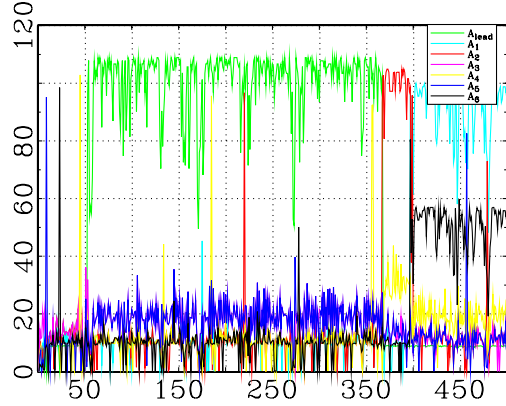


(c) Social structure in the final period ($t = 500$).

Figure 6: \uparrow recency: $\mu = 0.2$, $\lambda = 1$, $\phi_f = \phi_l = 0$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$, experience-based counterfactual estimation. Dominant leaders emerge despite the short memory but individual leaders last only a short period before someone else emerges. In the final period, Agent #77 leads most of the population but has been dominant as a leader for only 15 periods. The structure under 77 has a poor Δ -efficiency score.

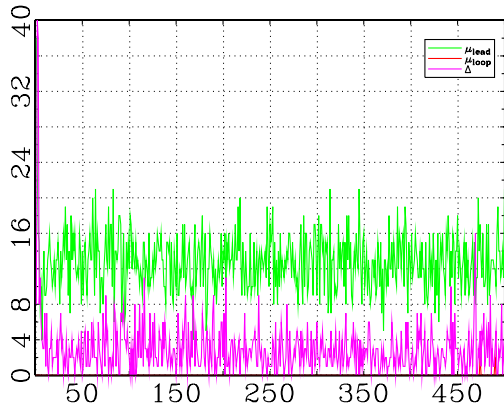


(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.




(b) Time-series of performance measures as experience by Agent #74.

Figure 7: \uparrow recency: $\mu = 0.2$, $\lambda = 1$, $\phi_f = \phi_l = 0$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$, expectations-based counterfactual estimation. Dominant leaders emerge. Despite the short memory, the dominant leader can retain the position for as long as she continues leads. Transitions between dominant leaders are quick and are the consequence of the dominant leader following another agent.




(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.




(b) Social structure in the final period ($t = 500$).

Figure 8: \downarrow Independence: $\mu = 0.2$, $\lambda = 1/3$, $\phi_f = \phi_l = 1$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$. A dominant leader emerges but in the presence of a persistent population of additional leaders that is larger when compared to the base simulation. Agent 24 is the dominate leader but not uniquely non-trivially.



(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.



(b) Social structure in the final period ($t = 500$).

Figure 9: \downarrow Follow Counterfactual: $\mu = 0.2$, $\lambda = 1/3$, $\phi_f = \phi_l = 1$, $\kappa_f = \kappa_l = 0$, $\delta_f = 0.6$, $\delta_l = 1$. No clear dominant leader emerges.

higher performing contacts means that the alternative to leading increasingly reflects the performance of just the superior performing contact(s). The larger population of leaders seen in Figure 8a is a consequence of the increased probability of leading among the distant followers of the leader, agent #24.

3.1.6 Counterfactuals

Decreasing δ leads agents to decrease the weight on estimated counterfactual payoff possibilities in the performance measure calculations. Actual experiences are given full weight. Hypothetical payoffs attributed to other actions are discounted. The failure to give weight to counterfactuals results in a loss of much of the information about the events transpiring around an agent. Agents fail to see the event to which they should respond and as a consequence, a dominant leader fails to emerge. With $\lambda = 1$, payoffs remain low regardless of strategy so that actions remain random throughout the simulation. With increasing independence, a social structure closer to autarky emerges as depicted in Figure 9.

3.2 Leadership Characteristics

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 randomly generated latent social network 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 10, those with a greater number of incoming links have a heightened probability for emerging as a leader. The distributions

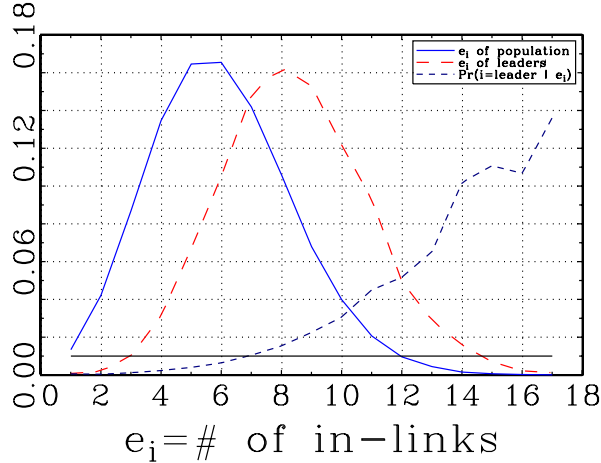



Figure 10: Distribution of e_i among the general agent population (solid blue) and among dominant leaders (long-dashed red). The conditional probability of leading (short-dashed black) increases in e_i .

and probabilities depicted in Figure 10 are generated from 5,000 simulation iterations. The solid (blue) curve is the distribution over the entire population of 500,000 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 .

The simulations used to generate the distributions in Figure 10 employ the baseline parameters with the modification that $A_0(1) = 10$. This establishes a prior belief that increases the probability of leading in the first period to 80%. The feature of interest in the baseline simulations is that a dominant leader typically emerges during the first 50-100 periods of the simulation and, once established, remains as leader throughout the simulation. Relative to the entire population, the distribution of e_i is shifted right towards higher values of e_i , with a sample mean $\bar{e}_i = 8.4$ for leaders relative to just $\bar{e}_i = 6$ for the entire population. Further, the conditional probability of emerging as a leader is increasing as e_i increases. The social advantage of possessing more incoming links heightens the probability of emerging as a leader, but this advantage does not completely dominate luck. Of the 6677 agents possessing only one incoming link, four emerged as the dominant leader. Agents possessing 17 or more incoming links have an estimated 13.6% probability of emerging as leader. This is considerably higher than the unconditional probability of 1%, but also far short of certainty.

A dominant leader is someone who both leads and attracts followers. The agent cannot be assured of the latter but does have control over the former. If x is the number of times an agent leads after q iterations of



(a) Baseline: All agents update probabilities according to the EWA. Mean earning is 31.88 with a std dev of 0.44.

(b) Treatment: Agent #1 always leads, all others update probabilities according to the EWA. For agents 2-100, mean earning is 32.69 with a std dev of 3.13. The per period average earning for Agent #1 is 22.32.

Figure 11: $\mu = 0.2$, $\lambda = 1$, $\phi_f = \phi_l = 1$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$, $T = 200$, 5,000 iterations. Distribution of per period average earnings across individuals. The behavior increases the probability that agent 1 emerges as the dominant leader of the simulation but lowers agent 1's expected earnings. The distribution of earning for agents #2 - #100 increases in variance relative to the base.

the simulation, then for a population of size n , $x \sim B(q, 1/n)$. For 5,000 iterations and $n = 100$, $E(x) = 50$ and $SD(x) = 7.04$. Consider n agents distributed over a symmetric social network of which $n - 1$ are homogeneous in their use of the EWA process to guide actions. Agent #1 deviates, choosing to always lead. Agent #1 emerges as the dominant leader in 1,102 of 5,000 iterations or just over 22% of the time. With agent #1 over-represented in the number of times as the dominant leader, the distribution of the number of observations as the dominant leader shifts down for agents #2 through #100, The narrow distribution is centered at 0.0079, very close to the $(1 - 0.21)/99$ likelihood for each agent #2 through #100 to emerge as dominant leader. Figure 11 presents the distribution of the per period average earning by each agent. Here, again, agent #1 stands out, earning considerably less than the remainder of the population. The foregone income from failing to follow when the strategy dominates exceeds the benefit gained from the increased probability of leading.

At the baseline parameters it is clearly not in the interest of an individual agent i to adopt a lead-only strategy. The higher probability of i emerging as the dominant leader presents the possibility that agent i may be better served by a strategy that deviates by persisting to lead until the performance of following dominates leading by a higher threshold than the EWA demands. The agent is effectively behaving strategically in an attempt to shape the evolution of the social structure while it is still malleable. Such a strategy can be approximated by increasing an individual's $A_0(0)$, the starting value for the lead performance

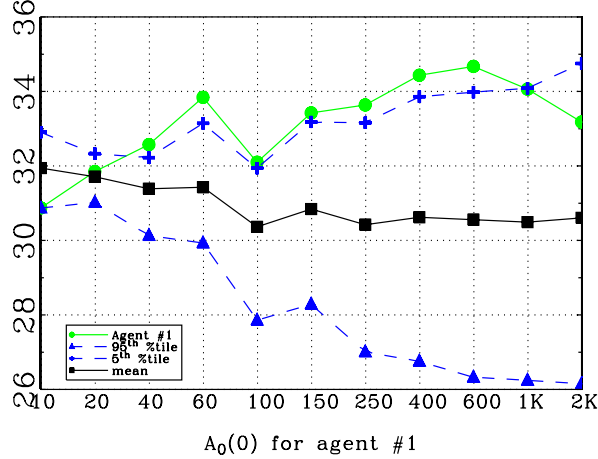


Figure 12: $\mu = 0.2$, $\lambda = 1$, $\phi_f = \phi_l = 1$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$, $T = 200$, 1,000 iterations. Average earning of agent #1's relative to the population as agent #1 increases persistence in leading. Population mean average earning and 5% bands shown.

measure. With $A_0(0) > A_d(0) = 0$, the agents start with a prior that is biased towards leading. For agents $j \in \{2, \dots, 100\}$, $A_{0,j}(0) = 10$. Seen in Figure 12 is the rise and then fall in agent #1's average earning relative to that of the population as $A_{0,1}(0)$ increases. With $A_{0,1}(0) = 10$, agent #1's average earnings is near the lower tail of the distribution. Agent #1's average earnings increases as she more persistently pursues the lead strategy, moving into the upper 5% tail of the average earnings distribution. It starts declining again at $A_{0,1}(0) = 1000$, though it remains above the average earnings with $A_{0,1}(0) = 10$.

Each agent has a position in the emerging social structure. An agent who deviates from taking her individual best position in the structure in order to lead lowers her own payoff and affects the entire population. Particularly affected are those whose distance minimizing route to the leader is increased due to the absence of the deviating agent from the tree. The deviant behavior also impacts who emerges as the dominant leader. Failing to adapt to another agent's observable success changes the probability of that agent emerging as the dominant leader. Overall, their deviation behavior slows the evolution towards an emerging structure. Figure 13 reveals that more time is required for a dominant leader to emerge when there are a subset of agents employing a lead-only strategy rather than adjusting behavior according to the EWA process.

3.3 With condition $B < 0$

A decrease in $(m-1)\frac{a_l}{a_T}$ 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 $\frac{a_l}{a_T}$ increases the relative premium reward to leading. Both

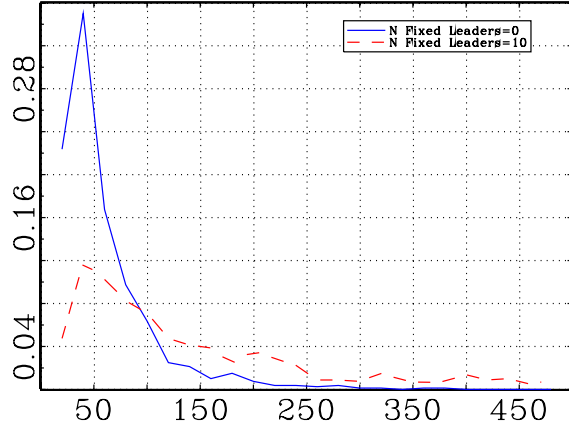


Figure 13: $\mu = 0.5$, $\lambda = 1$, $\phi_f = \phi_l = 1$, $\kappa_f = \kappa_l = 0$, $\delta_f = \delta_l = 1$, 800 iterations. Observed frequency where the horizontal axis indicates the number of periods before the first occurrence of a follower tree consisting of more than $n/2 = 50$ agents. The solid (blue) distribution is for simulations in which all agents employ the EWA while the dashed (red) distribution is based on simulations in which 10 of the 100 agents employ a lead-only strategy.

changes increase the incentive to lead, and for $B < 0$ the incentive is enough to induce the most distant followers of the dominant leader into leading.


In simulation, initially the full follower tree forms under the dominant leader. Once the tree approaches Δ -efficiency, the most distant followers realize the inferiority of following and begin to abandon the hierarchy. As the most distant followers abandon the tree, new agents find themselves to be the most distant followers and with diminished payoff. As the payoff to remaining a member of the tree of followers diminishes with the size of the follower population, so does the payoff to matching the dominant leader's selected option. From Goldbaum (2013), for $B < 0$, the equilibrium size of the follower population is the non-zero integer value just above or just below

$$n_h = \frac{1}{1 - (m - 1) \frac{a_J}{a_T}}. \quad (9)$$

Based on $B = -0.09$ producing $n_h = 10$, Figure 14 is typical of the $B < 0$ setting. The evolution is the same regardless of whether it is low m or low a_J/a_T that produces the $B < 0$.


3.4 Contact Selection

Introduced to the model is the ability to discard existing links to establish new links. This is implemented using a threshold \underline{w} such that for $w_{i,t}^j < \underline{w}$ agent i drops agent j from her list of contacts and randomly selects a new contact from the population.



(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.

(b) Time-series of performance measures as experience by Agent #100.



(c) Social structure in the final period (t = 500).

Figure 14: $\mu = 0.2$, $\lambda = 1$, $\phi_f = \phi_l = 0.98$, $\kappa_f = \kappa_l = 0.8$, $\delta_f = \delta_l = 1$, $T = 500$, $a_T = 1$, $a_J = 0.9/11 = 0.08\bar{1}$. $B < 0$ so that the equilibrium structure includes a single dominant leader and a pool of leaders without followers. Initially, the entire population follows agent #10 to form the Δ -efficient tree. Agents abandon the tree from the most distant positions as evidence accumulated to reveal the inferiority of being the most distant follower of a large tree.

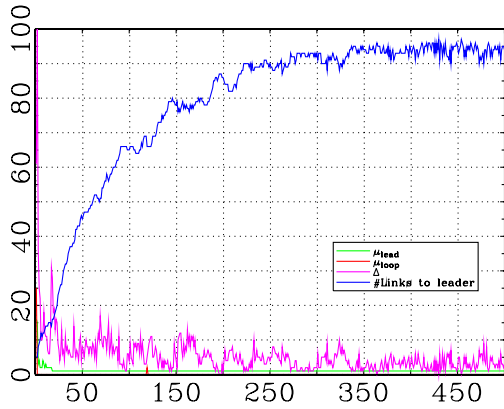
Allow individuals in the population to replace low performing contacts with new contacts and the hierarchy collapses the vertical tree structure as followers find more direct imitation paths to the leader. Eventually a completely horizontal structure emerges with the entire follower population directly imitating the dominant leader. The payoff to the leader is unchanged. The same is true for those who were most distant from the dominant leader according to the initial potential links structure. Those agents originally at intermediate distances from the leader see their payoff diminished by the collapse of the vertical structure as they lose their timing advantage over others in the population.

The simulation producing Figure 15 allows individuals to drop the lowest performing contacts with $w_{i,t}^j < \underline{w}_i$ with $\underline{w}_i = .5(1 - \theta_i)/d$. Under this criterion, the contact is dropped if he or she substantially underperforms with a weight that is only 50% of the average weight. By the final period of the simulation, nearly the entire population has developed a direct link to Agent #39, the leader in the population. By the time Agent #100 finds a direct link to the leader, around period 250, the premium earned for that direct link has diminished to near zero.


4 Conclusions

When the consumers are more concerned with the social phenomenon of a product than with the product itself, a choice leader plays a vital role in coordinating the population's adoption on a single product. Individuals have an interest in their position in the social structure when leading the population earns a premium reward. Despite the inherent inequality in the outcome, simulations indicate that a population can organize itself into the equilibrium social structure that benefits all when employing an appropriate mechanism of adjustment. The self-serving but myopic adaptive behavior produced by the Experience Weighted Attractor generates the needed adjustments in individual behavior that places the larger population into the role of follower behind a single emergent leader. That the emergence takes place with individuals unaware of the larger social structure addresses the Kirman et al. (2007b) criticism that endogenous network models tend to require unreasonably high knowledge regarding the full social network.


The failure in the population to organize under low Intensity of Choice conditions or when agents down-weight the performance signal of untried strategies points to the key elements necessary for structure to emerge. Outcomes early in the simulation are primarily driven by random events. In the absence of any accommodation to these early random events, they become transitory, subsumed by different randomized outcomes in the following period. Emergence of order requires adaptation that transforms early transitory



(a) Time-series of population characteristics. μ_{lead} = number who lead; μ_{loop} = number without a link to a leader; Δ = Δ -efficiency score.



(b) Time-series of performance measures as experience by Agent #100.



(c) Social structure in the final period ($t = 500$).

Figure 15: $\mu = 0.2$, $\lambda = 1$, $\phi_f = \phi_l = 0.98$, $\kappa_f = \kappa_l = 0.8$, $\delta_f = \delta_l = 1$, $T = 500$, $w_i = .5(1 - \theta_{i,t})/d$. With the ability to unilaterally and costlessly establish new potential imitation links, agents eventually find a direct link to the dominant leader.

outcomes into permanent components of a social structure. When agent i initially leads and is lucky in selecting an option that is popular this attracts the attention of those who can observe agent i increasing their likelihood of imitating i and thereby increasing the likelihood of a positive outcome for i in subsequent periods. Agent i 's success builds over time as the social structure adapts to, and thereby reinforces, her success.

The backward-looking adaptive accommodation by the agent to her own evolving environment governed by the EWA precludes the type of forward-looking strategic behavior an agent might attempt in order to influence the formation of the social structure to her own advantage. Personally advantageous enacted influence on the emergent social structure through forward looking strategic behavior is practiced to the detriment of the larger population. A preponderance of strategic actors delays or prevents the emergence of the equilibrium social structure.

References

- Acemoglu, D., Bimpikis, K., Ozdaglar, A. 2010. Dynamics of information exchange in endogenous social networks. Working Paper 16410, National Bureau of Economic Research.
- Acemoglu, D., Como, G., Fagnani, F., Ozdaglar, A. 2013. Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research*, 38, 1–27.
- Bala, V., Goyal, S. 2000. A noncooperative model of network formation. *Econometrica*, 68, 1181–1229.
- Blume, L. E., Durlauf, S. N. 2001. The interactions-based approach to socioeconomic behavior. In S. N. Durlauf, H. P. Young eds. *Social Dynamics*, Cambridge MA, MIT Press, 15–44.
- Branch, W., McGough, B. 2008. Replicator dynamics in a cobweb model with rationally heterogeneous expectations. *Journal of Economic Behavior & Organization*, 65, 224–244.
- Brock, W. A., Durlauf, S. N. 2001. Discrete choice with social interactions. *The Review of Economic Studies*, 68, 235–260.
- Camerer, C., Ho, T. H. 1999. Experience-weighted attraction learning in normal form games. *Econometrica*, 67, 827–874.
- Chang, M. H., Harrington, J. E. 2007. Innovators, imitators, and the evolving architecture of problem-solving networks. *Organization Science*, 18, 648–666.
- Dutta, B., Jackson, M. O. 2000. The stability and efficiency of directed communication networks. *Review of Economic Design*, 5, 251 – 272.
- Ellison, G., Fudenberg, D. 1995. Word-of-mouth communication and social-learning. *Quarterly Journal of Economics*, 110, 93–125.
- Goldbaum, D. 2013. Follow the leader: Equilibrium analysis of conformity and influence on a social network. SSRN Working Paper 1566611.
- Hausman, J., McFadden, D. 1984. Specification tests for the multinomial logit model. *Econometrica: Journal of the Econometric Society*, 1219–1240.
- Jackson, M. O., Watts, A. 2002. The evolution of social and economic networks. *Journal of Economic Theory*, 106, 265–295.

- Katz, M. L., Shapiro, C. 1985. Network externalities, competition, and compatibility. *American Economic Review*, 75, 424–440.
- Kirman, A., Markose, S., Giansante, S., Pin, P. 2007a. Marginal contribution, reciprocity and equity in segregated groups: Bounded rationality and self-organization in social networks. *Journal of Economic Dynamics & Control*, 31, 2085–2107.
- Kirman, A., Ricciotti, R. F., Topol, R. L. 2007b. Bubbles in foreign exchange markets. *Macroeconomic Dynamics*, 11, 102–123.
- Koenig, M. 2012. The formation of networks with local spillovers and limited observability. Discussion Paper No. 11-004, SIEPR.
- Schelling, T. C. 1971. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1, 143–186.
- Schelling, T. C. 1973. Hockey helmets, concealed weapons, and daylight savings - study of binary choices with externalities. *Journal of Conflict Resolution*, 17, 381–428.
- Watts, A. 2001. A dynamic model of network formation. *Games and Economic Behavior*, 34, 331–341.