A new approach to modeling social institutions using artificial neural networks

Simon Angus, Andrew Bell, Paul Hooper, Elizabeth Mullane, and Rafal Raciborski

August 15, 2007

1   Summary

Optimization models in biology and the social sciences assume that agents make decisions that maximize their payoffs under particular ecological, technological, and social constraints. While the assumption that the environment of decision-making is determined by factors outside the system often simplifies modeling, the fact that humans have the capacity to create and pass on institutions that markedly alter this environment greatly complicates attempts to model the evolution of human social behavior through time. While several attempts have been made to model the co-evolution of social institutions and behavior in fairly specific problem settings using analytic game theory and simulation (reviewed briefly below), we were interested in developing a new modeling framework which would allow the modeler to ask more general questions about the creation, maintenance, and evolution of social institutions abstracted from the details of a particular problem setting. This framework incorporates the use of artificial neural networks to represent the interrelated effects of behavior, institutions, and exogenous environment on individual payoffs, as well as an explicit network structure by which behavioral and institutional traits may move throughout a population. We present the first-generation results of this project, which investigate the impact of the structure of the informational network on the emergence and stability of individually and socially beneficial institutions.

2 Past approaches to social dilemmas and institutional evolution

We define institutions as artifices that are created or adopted by individuals within a system that impact the costs and benefits of a given set of behaviors, either by modifying expectations about how others in the system will behave (e.g. conventions), or by modifying the raw payoffs associated with a decision via reward or punishment (e.g. social norms, laws). Because there is no demand for institutions where institutionally unmediated self-interested decisions lead to good individual and social outcomes, most interest in the evolution of institutions focuses on social dilemmas—cases such as cooperation in the prisoner's dilemma, common pool resource use, and public goods provisioning—where myopic individual decision making is incompatible with achieving the social optimum (Gilbert 2006).

The evolution of cooperation among individuals has been tackled from a large number of perspectives. The existing economic literature approaches the study of norms through the calculus of a rational individual. Starting with Olson (1965), the underlying assumption is that agents will contribute to the common good up to the point where the marginal cost of cooperation equals its marginal benefit. Hardin's (1968) seminal paper on the "The Tragedy of the Commons" extended this logic to the problem of the overexploitation of common pool resources. These models assume that agents are perfectly informed about their and others' utility function and payoff structures (e.g. Ostrom, Gardner, & Walker 1992, Ostrom et al. 1999, Budescu, Rapoport, & Sulaiman 1995, Casari & Plott 2003), although as further studies show, this assumption is robust to relaxation (see, for example, Mookherjee & Sopher 1994, and Oechssler & Schipper 2003 for games of minimal information, and Apesteguia 2006 for games of incomplete information). The pessimistic conclusion is that cooperation is only possible in small groups or with an external enforcer.

Although analytically tractable, game theory models suffer from limitations with respect to the study of social norms and other institutions. First, most formal models provide only restricted means of accommodating institutional change. Even though a number of more recent studies introduce communication among players or between-round bargaining (e.g. Baland & Platteau 1996, Ostrom 1999, McCay 2002), agents negotiate only with respect to a specific institutional environment and change is represented by a new payoff structure instead of a new set of norms per se (see also Young 1993, Skyrms 1996, Boyd & Richerson 1996, Bowles 2001, and Henrich 2004 for formal evolutionary theories of institutions). The formal framework is not yet flexile enough to generalize the evolutionary dynamic to more abstract settings.

Evolutionary game theoretic models depart from the assumption of prescient rationality and investigate the behavioral change under given a process of selection or imitation (c.f. Lewontin 1961, Oster & Wilson 1978, and Maynard Smith 1982). Axelrod (1984, 1986) presents an evolutionary model of norms that relies on the elimination of inferior

and survival of beneficial strategies. Players retain the most effective strategies, and new strategies can be introduced by random mutations of old ones. Instead of converging to an equilibrium, the society often follows diverse paths of behavior.

3   A new approach

We were interested in developing a framework for modeling the co-evolution of individual behavior and social institutions utilizing artificial neural networks (ANNs). This framework has three appealing properties: it is at once novel, general, and generative.

Neural networks serve as multi-input, multi-output processing devices that can either be specified in a top-down manner, or trained via feedback processes to approximate unknown mapping functions. The sequential wiring of multiple internal layers, each composed of several nodes with non-linear activation functions, gives ANNs remarkable power in representing complex real-world functions (Krose & van der Smagt 1996). ANNs have been used to examine numerous problems in the physical and social sciences, including identifying specific oligonucelotides in the human genome, predicting fog, and determining the ration of browning in potato chips, to name just a few current applications (Liu et al. 2007, Fabbian, de Dear and Lellyett 2007, Serpen & Gokmen 2007). They have not, however, yet been systematically applied to understanding the adoption and evolution of social norms within a group.

The flexibility of the size and composition of the neural network make it a potentially powerful framework for analyzing general patterns in the development or collapse of social cohesion. We were interested in being able to ask questions about the role of learning rules and agent memory in promoting or attenuating the development of efficient social institutions without specifying, and therefore being restricted to, a particular problem setting (e.g. the public goods game or the repeated prisoner's dilemma). Like an ANN, each game theoretic problem setting specifies a mapping between agent behavior and payoffs, often in combinatorial, non-linear ways. Substituting the ANN for the specific payoff structure therefore allows for a generative approach, by which the experimenter may produce a class of game-theory-esque mapping functions, rather than being limited to addressing the specifics of one alone.

We imagine then that (exogenous) environmental and (endogenous) institutional effects are both realized through ANN processing functions. Behaviors may ultimately have positive, negative, or neutral effects on one's payoffs, effects which may be conditional on other actions undertaken by that same player or other players. Including an individual's behavior as well as an index (such as the mean) of aggregate group behavior as inputs to the payoff function allows us to implement positive or negative externalities that individual decision-making has on other players.

4 A first-generation model

Norms and environmental constraints are built into the model as a two-layer neural network with n outer nodes and m inner nodes. The n outer nodes represent n independent traits, or behaviors, of the agents. The m inner nodes can be conceptualized as m independent relationships among these behaviors. Thus, the size of m relative to n can be thought of as describing the complexity of interactions among behaviors. The first layer of connections (from n nodes to m nodes) represents social norms—collective understanding of how behaviors relate to each other and the environment—and changes over the course of each experiment. The second layer of connections (from m nodes to n nodes) represents the environmental constraints imposed upon the different behaviors, whose values are randomly assigned and which remain fixed over the course of the experiment. The society consists of a set of q agents, which are related through the q×q network structure matrix S, an adjacency matrix of connections among agents.

Each agent is randomly assigned values scaling between 0 and 1 for each trait, and values scaling between -1 and 1 for each connection in the first (norms) layer of the neural network. The actual norms layer is calculated as the average across all agents for the m×n set of connections. The environmental constraints are also assigned values scaling between -1 and 1 at the beginning of the experiment. The matrix S is assigned values of 0 and 1 randomly, with the density of 1s (connections) assigned by the parameter f ranging from 0 to 1.

Agent payoffs are calculated by inputting the set of traits to the neural network, and calculating the output from the two layers. If E is the adjacency matrix describing the environmental constraints, W is the adjacency matrix describing the averaged norms across agents, and b is the vector of traits for agent i, then the payoff vector p is given by:

$$p = W_{n \times m} * E_{m \times n} * b_{n \times 1}$$

The overall payoff for an agent is simply the sum of the payoff vector p. Agents improve their payoffs over time by looking back through a "memory" of length k of previous payoffs. The value of the lowest payoff in memory is subtracted from all payoffs, and the resulting values converted to a cumulative sum vector, normalized to range from 0 to 1. The distance between adjacent values i-1 and i in this vector represents the probability that i is chosen as the best payoff, so that larger payoffs have a higher probability of being selected. Once a best payoff is selected, agents choose with equal probability to update either their behavior b or norms W with the associated b or W from that time period. In the case of W, the updated W will actually be a weighted average of all agent W values from that time period for agents connected to the current agent (as described in S).

After updating, all values for the agent b vector and W matrix are mutated by a randomly selected value within the ranges δ_b and δ_W. A constraint in this model is that the sum of the b vector (effort put into behaviors) and the sum of the absolute values of the W matrix (effort, or conviction, for different ideas) are chosen to be conserved. Thus, after

Parameter	Value
Outer Nodes n	3
Inner Nodes m	5
Agents q	20
δ_b	0.0667
δ_W	0.0133

Table 1: Parameter Values for Experiments 1 and 2

updating and mutating agent traits, both b and W are renormalized to retain the same summed value as in each previous timestep. Thus, adding strength to any one connection or behavior indirectly takes strength away from all other connections or behaviors.

This process is repeated for all agents in the society, for all timesteps in the simulation.

5 Results

5.1 Effect of Information Structure

For a small number of initial random seeds, we varied the network density parameter f from 0 (no connections) up to a value of 0.9 (densely connected network). We performed two sets of experiments—in Experiment 1, the values of W and E were allowed to scale from -1 to 1; in Experiment 2, the values of W and E were allowed to scale only from 0 to 1 (Table 1).

For a typical run in Experiment 1, the average payoff for all agents rapidly increases as the update and mutate mechanism quickly allows each agent to improve the allocation of their efforts (Figure 1 A). After an initial period of rapid growth, the growth slows, and trends become non-monotonic. It is interesting to see that as f increases from 0, the path of the average payoff appears to become less smooth for small non-zero values of f. However, for larger values of f, the average payoff is much smoother and lower overall. As f is increased, the rate at which the traits b settle out increases dramatically (Figure 1 B), so that for high values of f the variance in b across agents is at a minimum from very early on—the agents are not exploring the trait space. The variance will however not drop to zero, as there is some variance in the total 'effort' available to each agent, which is conserved throughout the experiment. The average variance in W across agents is also lower for higher values of f (Figure 1 C). Here, the width of the path is largely a function of the δ_W parameter; what is striking is the lack of large spikes in average variance for larger values of f, seen in the other runs as major shifts in 'norms' occur.

Much is the same for a typical run in Experiment 2, with one striking exception (Figure 2 A). Here, the shapes of all paths are similar to Experiment 1, except that now larger values of f lead to higher average payoffs. A more subtle difference in the results is that
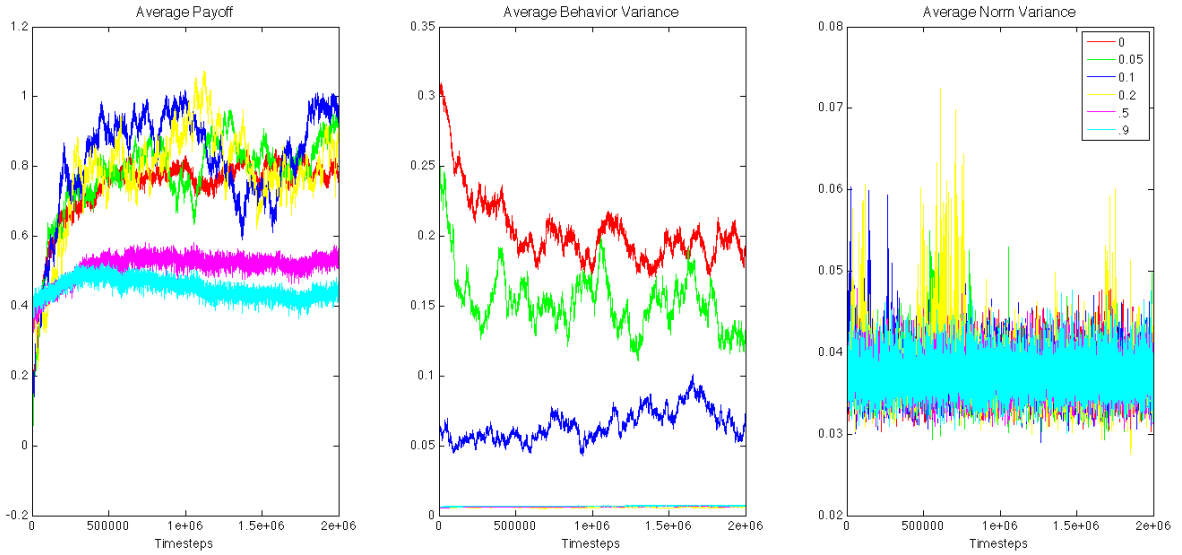
Figure 1: A typical model run with W and E scaling from -1 to 1

the average variance in traits b appears to be very slowly tracking upward from initially low values for runs with larger f.

6   Discussion

6.1   Information Structure

One clear and easily understood effect of information structure is that trends in average payoff are smoother with higher values of f. The averaging effect of information exchange among agents makes their own traits and norms more homogeneous and less vulnerable to large swings.

The more interesting effect is the clear difference in the value of information between cases where W and E scale between -1 and 1 (Experiment 1), and where they only scale from 0 to 1 (Experiment 2). One possible explanation for this is that when negative values are allowed, the averaging of W over large numbers of agents may give results for W that differ in sign from that of the true optimal value. A negative sign in W when the corresponding link in E is positive will give a negative payoff, and vice versa. In cases where negative values are not allowed, however, this issue does not arise—the averaging effect simply smoothes out the norm values, and perhaps leads to better overall outcomes.

There is a clear slow trend in behavior variance for high values of f in Experiment 2. Possibly, the smoother, more coherent averaging effect in Experiment 2 provides a better
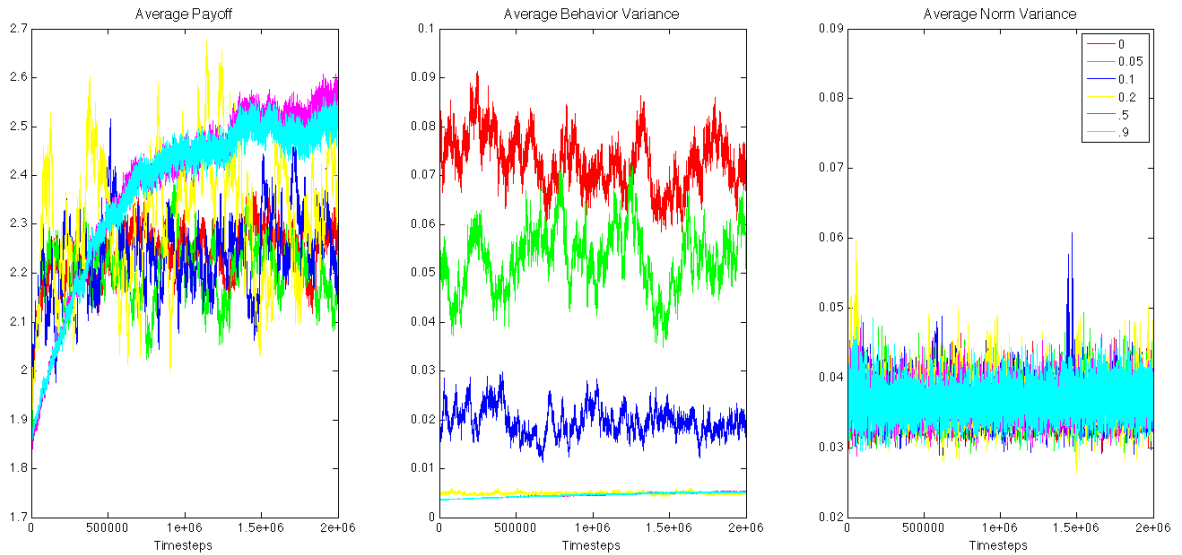
Figure 2: A typical model run with W and E scaling from 0 to 1

environment for agents to explore different behaviors when information is high than under the same conditions in Experiment 1.

Also, though the entire f space is not explored, it seems as though there exists some kind of information 'percolation threshold' separating two behavioral regimes, a value less than 0.5 and greater than 0.2. Between 0 and this threshold value, the sharing of information only among pockets of agents appears to allow large swings in average payoff across the society, and may be leading to locally dominant behavior and norm strategies.

7   Future work

We plan to develop a second generation of models utilizing the ANN framework that incorporates the following features: (1) institutional and environmental processing will be functionally separated, allowing for non-zero payoffs even in an institutionless 'state of nature'; (2) average population behaviors will be included as inputs to both institutional and environmental ANNs, allowing for the explicit representation of positive and negative externalities; and (3) we will allow investment in institutions to impose direct costs on agents, such that one must sacrifice wealth to carry normative beliefs or support institutional mechanisms of reward or punishment.

8   Acknowledgements

9 References

1. Apesteguia, J. (2006). Does Information Matters in the Commons? Experimental Evidence. Journal of Economic Behavior & Organization, 60, 55-69.

2. Axelrod, R. (1984). The evolution of cooperation. New York: Basic Books.

3. Axelrod, R. (1986). An Evolutionary Approach to Norms. American Political Science Review, 80(4), 1095-1111.

4. Axelrod, R. (1997). The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton, NJ: Princeton University Press.

5. Axelrod, R. An evolutionary approach to norms. Am. Pol. Sci. Rev. 80, 1095-1111 (1986).

6. Baland, J., & Platteau, J. P. (1996). Halting Degradation of Natural Resources, Is There a Role for Rural Communities? Oxford: FAO and Clarendon Press.

7. Bowles, S. (2001). Individual Interactions, Groups Conflicts, and the Evolution of Preferences. In S. Durlauf & P. Young (Eds.), Social Dynamics (pp. 155-190). Cambridge, MA: MIT Press.

8. Boyd, R. & Richerson, P.J. Culture and the Evolutionary Process (University of Chicago Press, Chicago, IL,1985).

9. Boyd, R., & Richerson, P. J. (1996). Why Culture is Common, but Cultural Evolution is Rare. In W. G. Runciman, J. M.

10. Smith & R. I. M. Dunbar (Eds.), Evolution of Social Behavior Patterns in Primates and Man (pp. 77-93). Oxford: Oxford University Press.

11. Budescu, D. V., Rapoport, A., & Suleiman, R. (1995). Common Pool Resource Dilemmas Under Uncertainty: Qualitative Tests of Equilibrium Solutions. Games and Economic Behavior, 10, 171-201.

12. Cannibal, G.L. & Winnard, G.M. Managing the tragedy: an inter-disciplinary model for managing the environment as interacting chaotic hierarchy. Futures. 33: 2,147-160 (2001).

13. Crdenas, J-C. & Ostrom, E. What do People Bring into the Game? Experiments in the field about cooperation in the commons. Agricultural Systems. 82: 3, 307-326 (2004).

14. Casari, M., & Plott, C. H. (2003). Decentralized Management of Common Property Resources: Experiments with a Centuries-Old Institution. Journal of Economic Behavior & Organization, 51, 217-247.

15. Durrett, R., & Levin, S. A. (2005). Can Stable Social Groups be Maintained by Homophilous Imitation Alone? Journal of Economic Behavior & Organization, 57, 267-286.

16. Esteban, J. & Ray, D. Collective action and the group size paradox. American political science review. 95: 3, 663-672 (2001).

17. Fabbian, D., de Dear, R. & Lellyett, S. Application of artificial neural network forecasts to predict fog at Canberra International Airport. Weather and forecasting. 22: 2, 372-381 (2007).

18. Faysse, N. Coping with the tragedy of the commons: Game structure and design of rules. Journal of economic surveys. 19: 2, 239-261 (2005).

19. Gilbert, M. Rationality in Collective Action. Philosophy of the social sciences. 36: 1, 3-17 (2006).

20. Grant, W.E,. & Thompson, P.B. Integrated ecological models: simulation of sociocultural constraints on ecological dynamics. Ecological modeling. 100: 1-3, 43-59 (1997).

21. Hardin, G. The Tragedy of the Commons. Science. 162, 1243-1248 (1968).

22. Henrich, J. (2004). Cultural Group Selection, Coevolutionary Processes and Large-Scale Cooperation. Journal of Economic Behavior & Organization, 53, 3-35.

23. Holley, R., & Liggett, T. M. (1975). Ergodic Theorems for Weakly Interacting Systems and the Voter Model. Annals of Probability, 3, 643-663.

24. Krose, B. & van der Smagt, P. Introduction to Neural Networks (University of Amsterdam, Amsterdam, 1996).

25. Lewontin, R. C. (1961). Evolution and the Theory of Games. Journal of Theoretical Biology, 1, 382-403.

26. Lipford, J.W. Group-size and the free-rider hypothesis - an examination of new evidence from churches. Public choice. 83: 3-4, 291-303 (1995).

27. Liu, C. et al. Genome-wide identification of specific oligonucleotides using artificial neural network and computational genomic analysis. BMC Bioinformatics. 8:164, (2007).

28. Macy, M. W., Kitts, J. A., & Flache, A. (2004). Polarization in Dynamic Networks: A Hopfield Model of Emergent Structure. In R. Breiger, K. Carley & P. Pattison (Eds.), Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers (pp. 162-173). Washington, DC: The National Academy Press.

29. Maynard Smith, J. (1982). Evolution and the Theory of Games. Cambridge: Cambridge University Press.

30. McCay, B. J. (2002). Emergence of Institutions for the Commons: Contexts, Situations, and Events. In E. Ostrom, T. Dietz,

31. N. Dolsak, P. C. Stern, S. Stonich & E. U. Weber (Eds.), The Drama of the Commons, Division of Behavioral and Social Sciences and Education. Washington, DC: National Academic Press.

32. Mookherjee, D., & Sopher, B. (1994). Learning Behavior in an Experimental Matching Pennies Game. Games and Economic Behavior, 7, 62-91.

33. Nakamaru, M., & Levin, S. A. (2004). Spread of Two Linked Social Norms on Complex Interaction Networks. Journal of Theoretical Biology 230, 57-64.

34. Oechssler, J., & Schipper, B. (2003). Can You Guess the Game You Are Playing? Games and Economic Behavior, 43, 137-152.

35. Olson, M. (1965). The Logic of Collective Action. Cambridge, MA: Harvard University Press.

36. Olson, M. The Logic of Collective Action (Harvard University Press, Cambridge, MA, 1965).

37. Oster, G. F., & Wilson, E. O. (1978). Caste and Ecology in the Social Insects. Princeton: Princeton University Press.

38. Ostrom, E. (1999). Coping with Tragedies of the Commons. Annual Review of Political Science, 2, 493-535.

39. Ostrom, E. Collective Action and the Evolution of Social Norms. Journal of Economic Perspectives. 13/4, 137-158 (2000).

40. Ostrom, E., Burger, J., Field, C. B., Norgaard, R. B., & Policansky, D. (1999). Revisiting the Commons: Local Lessons, Global Challenges. Science, 284, 278-282.

41. Ostrom, E., Gardner, R., & Walker, J. (1994). Rules, Games, and Common-Pool Resources. Ann Arbor, MI: University of Michigan Press.

42. Pandit, B.H. & Thapa, G.B. A Tragedy of Non-Timber Forest Resources in the Mountain Commons of Nepal. Environmental conservation. 30: 3, 283-292 (2003).

43. Poteete, A.R. & Ostrom, E. Heterogeneity, group size and collective action: The role of institutions in forest management. Development and change. 35: 3, 435-461 (2004).

44. Serpen, A. & Gokmen, V. Modeling of acrylamide formation and browning ratio in potato chips by artificial neural network. Molecular nutrition & food research. 51: 4, 383-389 (2007).

45. Singleton, S. Commons Problems, Collective Action and Efficiency: Past and Present Institutions of Governance in Pacific Northwest Salmon Fisheries. Journal of Theoretical Politics. 11: 3, 367-391 (1999).

46. Siren, A.H. Natural resources in indigenous peoples' land in Amazonia: A tragedy of the commons? International journal of sustainable development and world ecology. 13: 5, 363-374 (2006).

47. Skyrms, B. (1996). Evolution of the Social Contract. New York: Cambridge University Press.

48. Worden, L. & Levin, S.A. Evolutionary Escape from the Prisoner's Dilemma. Journal of theoretical biology. 245: 3, 411-422 (2007).

49. Young, H. P. (1993). The Evolution of Conventions. Econometrica, 61, 57-84.