Genotype or Phenotype?

The conflation of two concepts in evolutionary agent-based modeling.

Will Braynen, Simon Angus, Paul Dwyer, Mollie Poynton, Alejandro Balbin,
and Risi Kondor[1]

ABSTRACT. The two distinct concepts of genotype and phenotype are often conflated by Neo-Darwinian evolution theory, as well as by agent-based models that use evolutionary principles derived from that theory (e.g. in evolutionary game-theoretic models). In evolutionary game-theoretic models, the concept of genotype is analogous to an agent's strategy while the concept of phenotype is analogous to an agent's behavior. Strategy and behavior, however, are often treated as equivalent concepts. Using a spatialized prisoner's dilemma, we show this conflation is a non-trivial simplification that may lead to drastically different outcomes. Additionally, we discuss implications for using the evolutionary paradigm to model market interactions.

I. CONCEPTS: GENOTYPE AND PHENOTYPE

In biology and population genetics, genotype is the genetic constitution of an organism (as specified by its alleles) (Lewontin). The phenotype of an agent, by contrast, is a characteristic expression of the agent's genotype given its environment (Lewontin). Schematically,

genotype + environment + random-variation → phenotype.

To simplify and emphasize the relationship of interest here,

genotype + environment → phenotype.

Phenotype + environment, in turn, gives rise to fitness upon which (natural or social) selection operates. For example, if you have a genetic predisposition to Type 2 diabetes, you might never become diabetic. Whether or not a person genetically predisposed to Type 2 diabetes becomes diabetic (genotype) depends largely on environment. Even though a predisposition to Type 2 diabetes seems undesirable (poor fitness), a person possessing that genotypic predisposition might never be exposed to a sugar-heavy diet and so never develop the disease. Instead, the genotype would be hidden, and the successful phenotype would be attractive to potential mates (high fitness). Importantly,

---

[1] Will Braynen, Simon Angus, Paul Dwyer, Alejandro Balbin and Risi Kondor are doctoral candidates at the University of Arizona, the University of New South Wales, Texas A&M University, the National University of Columbia and Columbia University respectively. Mollie Poynton is an Assistant Professor at the University of Utah. This work was partially supported by the Santa Fe Institute through NSF Grant No. 0200500 entitled "A Broad Research Program in the Sciences of Complexity."

the genotype is not directly observable, and it is only phenotype, the expression, that is observable by others.

These concepts are also relevant when modeling social phenomena. When applying these evolutionary concepts to the modeling of social phenomena, genotype is analogous to an agent's strategy, and phenotype is analogous to the agent's behavior. For example, a stock broker's genotype (strategy) may be to buy shares of Microsoft, Inc. whenever it rains. However, a summer long dry spell may make it appear as though the stock broker never buys shares of Microsoft, Inc. His observable behavior (phenotype) gives no indication he would ever buy shares of Microsoft, Inc.

II. CRITICISM OF CURRENT DOGMA: HOW SOCIAL DARWINISM DIFFERS FROM BIOLOGICAL DARWINISM

Social evolution differs from biological evolution in that social evolutionary mechanisms can operate within a single generation (*intra*generationally), as opposed to operating only between generations (*inter*generationally). It's true that successful agents can pass on their strategy (their genotype) to their offspring – be the offspring children (if the agent is a parent) or franchises (if the agent is Starbuck's Coffee). This is intergenerational social evolution, similar to biological evolution. Social evolution is distinctive from biological evolution in that successful agents do not merely propagate strategy to offspring, but also serve as exemplars and role models. The behavior of successful agents is often imitated by other agents (whose strategy might be to partially or wholly imitate the behavior of successful agents).

A common problem in social modeling, in particular evolutionary game-theoretic models, is that agents are usually programmed only to copy the strategy (genotype) of other successful agents. In other words, agents are designed to imitate that which they could not possibly observe or re-construct from publicly accessible information. Only the behavior (phenotype) is observable. In the real world, *within a single generation*, we often copy each other's behavior, but not the contents of each other's heads (beliefs, desires, and strategy). Agents in a real economic market, for example, cannot imitate what's internal to another agent. Furthermore, copying the behavior of a successful agent may not always lead to success. For example, Canon might invest in making cheap color printers not because Canon expects to make money off printers, but because it hopes to make money from selling ink cartridges. (And conversely, a successful imitator might be successful in the short run not because she is using a good strategy, but because she merely got lucky and copied behavior resulting in high fitness.) In short, imitation of behavior and reconstruction of strategy happen at the level of the observable, not at the level of the internal structure , whereas computational economic models (e.g. agent-based models that use evolutionary game theory) currently model interactions as if imitation and learning proceeds at the internal, publicly inaccessible level of copying the strategy of other agents who are successful. This conflation of genotype and phenotype – of unobservable strategies and observable behavior – is a point of concern in agent-based modeling, as it yields very different results. We demonstrate this concern with a simulation.

III. THE MODEL

There is more than one way to distinguish between behavior and the strategy underlying that behavior. One way to draw the distinction is to differ the space of possible behaviors from the space of possible strategies. Consider the classic strategy of tit-for-tat as an example.[2] An agent playing tit-for-tat will cooperate on the very first round and thereafter will cooperate with another agent only if the agent cooperated on the previous round (and will defect if the other agent defected on the previous round). Essentially, tit-for-tat is a conditional strategy. If you observe a tit-for-tat player interact with a cooperator, you will never see the tit-for-tatter defect. To an outside observer, a tit-for-tat player will be indistinguishable from an unconditional cooperator.

Consider a world of only two possible actions – cooperation and defection – in which the strategy is conditioned on another agent‘s behavior in the previous round (and only on the previous round). There are eight possible strategies, conditioned on the following: whether the agent cooperates the very first time it meets someone new (two possible actions here), whether the agent cooperates if its neighbor cooperated on the previous round (two possible actions), and whether the agent cooperates if its neighbor defected on the previous round (two possible actions). In total, there are eight (2x2x2) possible conditional strategies.[3] But while there exist eight possible strategies, there exist only two possible behaviors: cooperation and defection.[4] So, behavior and strategy can come apart conceptually if the space of possible behaviors is a subset of the possible strategies.[5]

There is, however, a way to distinguish strategy from behavior without making the space of possible behaviors smaller than the space of possible strategies. The way we do this in our model is by making the strategy probabilistic. In other words, an agent's strategy is the probability with which she cooperates with others. Say, for example, that my probability of cooperation is 0.5. Every time I have to decide whether to cooperate with someone, I flip a fair coin. In the long-run, following this strategy, we would expect that I cooperate about half the time. If, however, you watch me interact with others for only ten rounds, it is very unlikely that I will cooperate for exactly five of those ten rounds. The more limited your sample of observations, the more likely it is that the behavior you observe deviates from the underlying probabilistic strategy. This is the model we implemented, precisely because it allows to leave the space of behaviors as large as the space of strategies while still distinguishing between them.

In our model, we used a spatial layout in which agents (situated in a lattice) interact with their neighbors.[6] We use probabilistic strategies (an agent's probability of

---

[2] See Axelrod (1984) for a study of tit-for-tat.

[3] A tit-for-tat player, for example, has a **C**-**C**-***D*** strategy because she **C**ooperates the first time she meets a new player, **C**ooperates if the other player cooperated on the previous round, and ***D***efects if the other player defected on the previous round.

[4] This assumes, of course, a rather simple-minded outside observer.

[5] Notice that if we allowed only two possible conditional strategies into the population of C-C-C and D-D-D, then this space of possible strategies would be the same as the space of possible behaviors C (cooperate) and D (defect).

[6] We chose this setup because of the precedent set it in the literature; see Nowak and May's seminal papers (1992, 1993), as well as Grim (1995, 1998).

cooperation) for the reasons outlined above. In our model, as in similar published models (see Nowak and May 1992, 1993, as well as Grim 1995, 1998), each agents sums the payoffs it receives from interacting with each of its neighbors. The payoffs from the interaction with neighbors are governed by Nowak and May's prisoner's dilemma payoffs.[7]

Our agent-based model, similar to other published models proceeds with the following steps:

1. Action decision: Each agent flips a biased coin to decide whether or not to cooperate. (The coin is biased by its strategy: by its probability of cooperation [0,1]. In the initial population, there is a normal random distribution of strategies across agents.)

2. Interactions with neighbors: The agent interacts with one of its neighbors (who also flips a biased coin to decide whether or not to cooperate). The payoff depends on whether the agents cooperate or defect and, in the case of unilateral cooperation, on who cooperated.

3. Payoffs from resulting interactions: Step (2) is repeated for each of the agent's neighbors and payoffs are summed from all interactions of an agent with its neighbors.

4. The evolutionary/learning step: Each agent imitates its most successful neighbor (but only if one of its neighbors did better, i.e. received a higher payoff, than the agent).

Because our goal is to see whether or not the distinction between strategy and behavior – between genotype and phenotype – matters, we created two models in which step four differs. In both models, interactions proceed as outlined above. In model A , agents imitate the most successful neighbor's *strategy*, copying its underlying probability of cooperation. Model A is the conflation model. However, in model B, agents imitate *behavior*, copying the observed probability of cooperation. Model B is the realistic model.

Consider a successful neighbor with an underlying strategy of 0.5 (genotype = 0.5), and whose observed behavior in the last round was to cooperate in 6/8 interactions (observed behavior or phenotype = 0.75). In model A (imitation of genotype), agents imitate the strategy (the genotype) of their most successful neighbor. Agents copy the *probability* (or probabilistic propensity) towards cooperation of another more successful agent. So in our example, an agent will use a coin biased 0.5 in favor of cooperation when deciding whether to cooperated with its neighbors on the next round. You can think

---

[7] In general, a two-player prisoner's dilemma is defined as follows: CD < DD < DC < CC, with CC > (CD+DC)/2, where CC is the payoff for mutual cooperation, DC for unilateral defection, DD for mutual defection, and CD for unilateral cooperation. For robustness, we used two different sets of prisoner's dilemma payoffs: <CD,DD,DC,CC> = <0,1,3,5> and <0,1,5,6>. The first set of payoffs <0,1,3,5> is the more traditional set used by Axelrod (1984), while <0,1,5,6> is a set of payoffs more conducive cooperation (Nowak and May,1992)

of this as our control case. In model B (imitation of phenotype), agents imitate the behavior (the phenotype) of their most successful neighbor. Agents copy the *frequency* with which another more successful agent was observed to cooperate. So in our example, an agent will use a coin biased 0.75 in favor of cooperation when deciding whether to cooperated with its neighbors on the next round (because its most successful neighbor was observed to have actually cooperated 0.75 of the time).[8]

IV. METHODS

The model was implemented in Wilensky's (1999) *NetLogo* modeling environment (Version 4.0 beta 3). A screenshot of the implementation user interface is given in Figure 1. The simulation took place within a 23 by 23 cell toroidal (i.e., doughnut-shaped) space. A series of 200 round experiments were conducted using both model A and model B, first with Nowak & May payoffs, then Axelrod payoffs. Neighborhood radius was varied from 1 to 4.

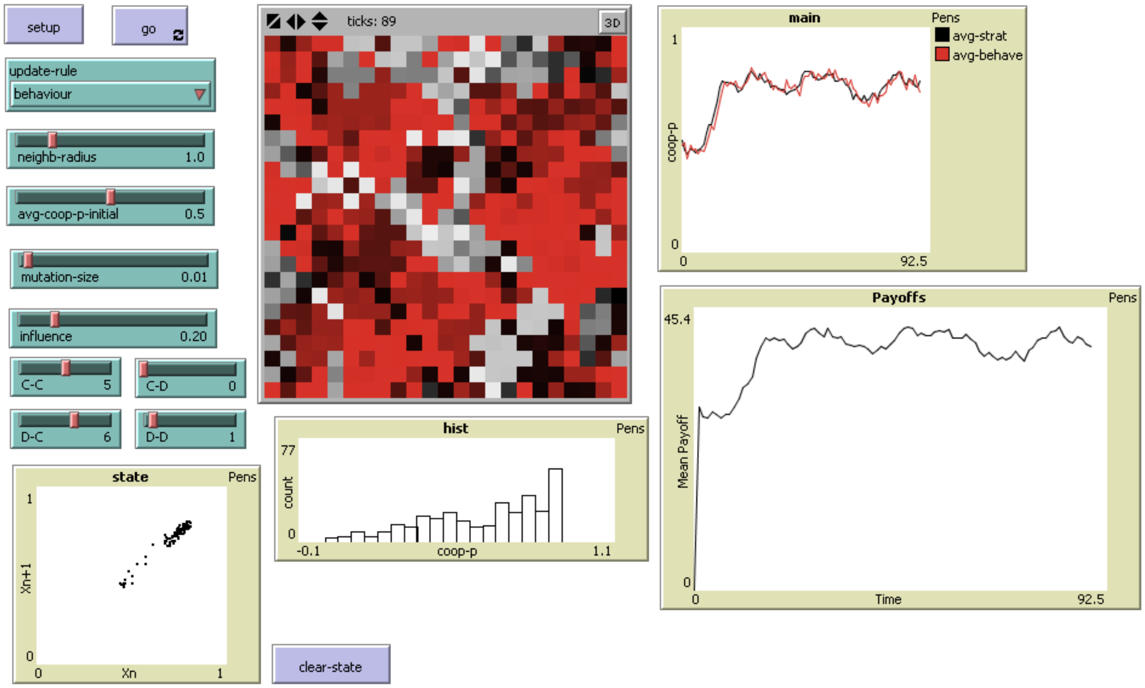

update-rule: behaviour
neighb-radius 1.0
avg-coop-p-initial 0.5
mutation-size 0.01
influence 0.20
C-C 5
C-D 0
D-C 6
D-D 1
ticks: 89

Figure 1: NetLogo Simulation User Interface

---

[8] Given a large enough sample, of course, the observed frequency of cooperation and the underlying probability of cooperation should converge to the same value. This, however, would wash away the distinction between strategy and behavior in our model. So keeping the sample size small is crucial in our model to maintaining the strategy/behavior (genotype/phenotype) distinction.

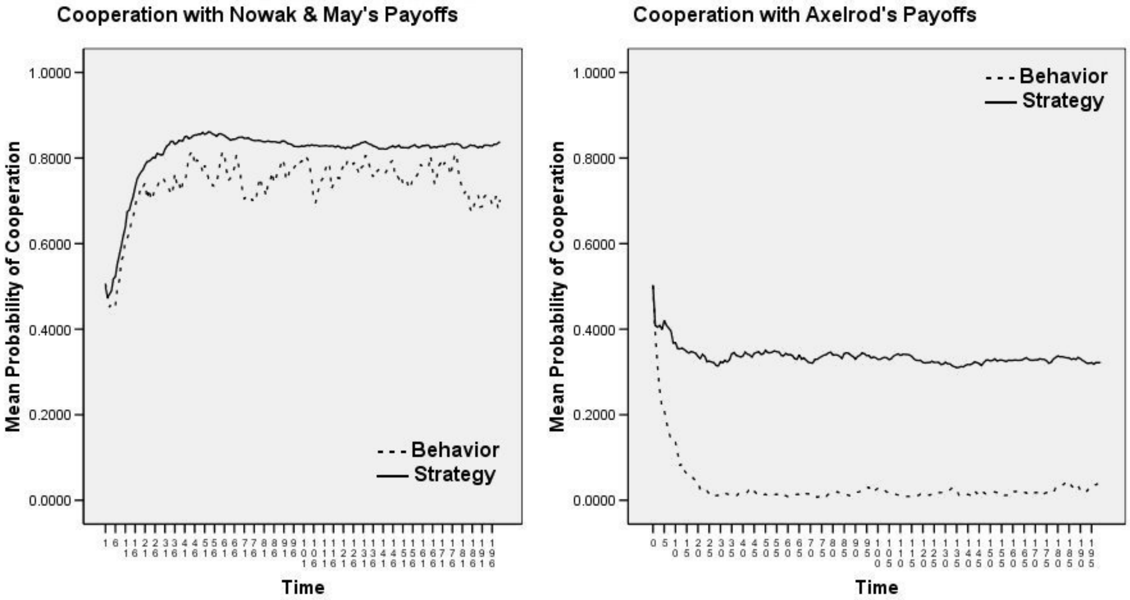

Figure 2: Robustness of Cooperation across Payoffs (One Neighbor)[9]

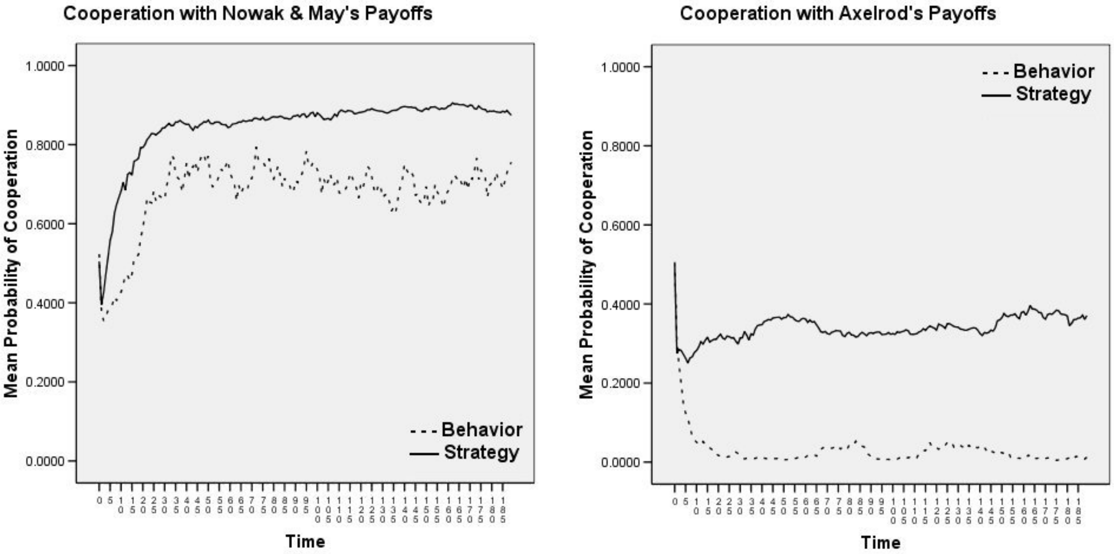

Figure 3: Robustness of Cooperation across Payoffs (Two Neighbors)

---

[9] For robustness, we used two different sets of prisoner's dilemma payoffs: <CD,DD,DC,CC> = <0,1,3,5> and <0,1,5,6>. The first set of payoffs <0,1,3,5> is the more traditional set used by Axelrod (1984), while <0,1,5,6> is a set of payoffs more conducive cooperation (Nowak and May,1992)

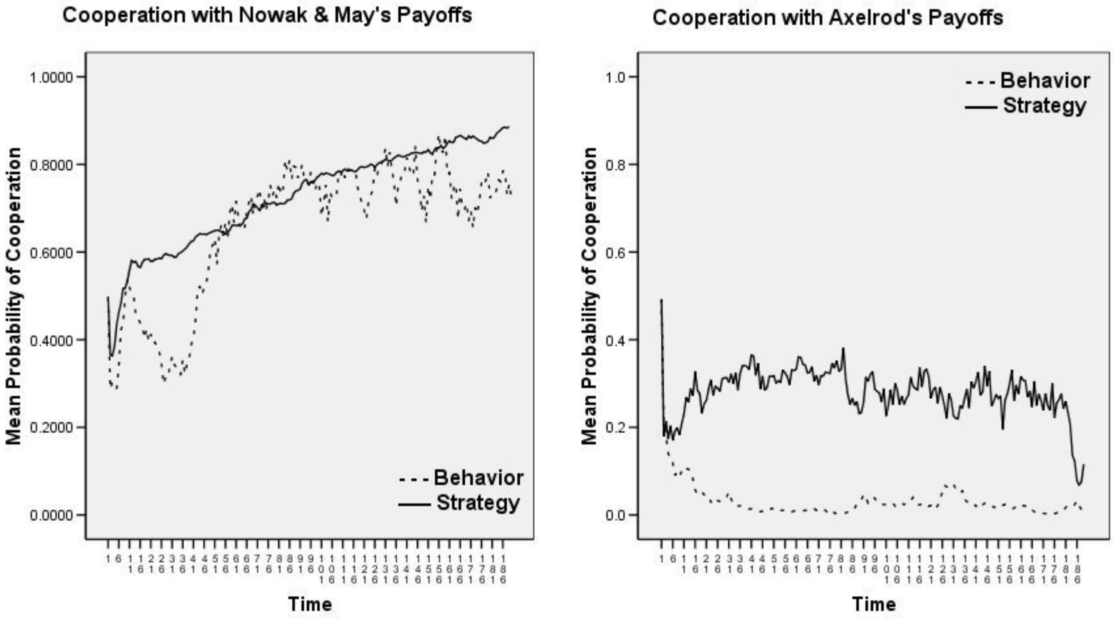

Figure 4: Robustness of Cooperation across Payoffs (Three Neighbors)

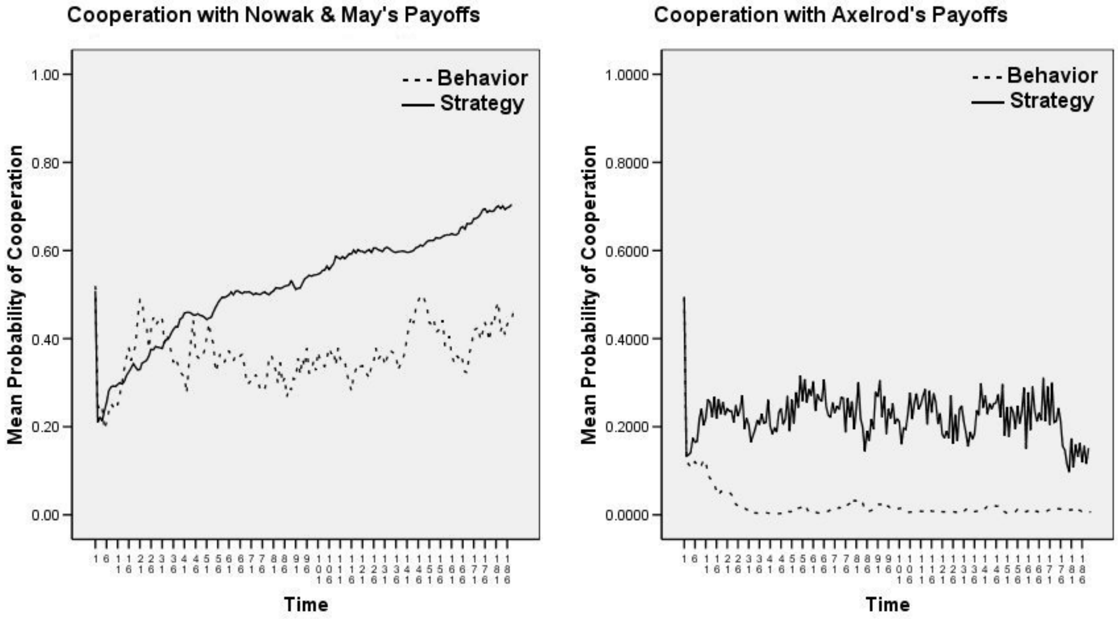

Figure 5: Robustness of Cooperation across Payoffs (Four Neighbors)

V. RESULTS

Figures 2 through 5 show the results according to model, payoff, and neighborhood radius. It is readily apparent that imitating the strategy (i.e., genotype) of one's more fit neighbors results in better outcomes than merely imitating their behavior. Genotype imitators are more cooperative and therefore will earn greater payoffs regardless of whether Nowak and May, or Axelrod payoff matrices are used. These results are robust across neighborhood radius values. The mean probability of cooperation that resulted from agents imitating the *strategy* of their more successful neighbors was consistently higher than the mean probability of cooperation that resulted when agents imitated *behavior*.. For example, with Axelrod payoffs of <CD,DD,DC,CC> = <0,1,3,5> and a neighborhood radius = 1 (a Von Neumann neighborhood of only four neighbors), agents evolve to have a probability of cooperation equal to 0.0 if behavior is imitated.[10] If, however, agents imitate the strategy of other agents instead of their behavior, the average probability of cooperation after 100 generations is as high as 0.4.[11]

VI. DISCUSSION

Both individual agent fitness and overall population dynamics may differ substantially in an evolutionary game theoretic model, depending on whether the learning agent copies the phenotype (behavior) or the genotype (strategy) of other agents. We used continuous strategies (probability of cooperation), but the same result would be expected with discrete conditional (reactive) strategies: e.g. tit-for-tat can appear as an unconditional cooperator to an outside observer. The dynamics differ because behavioral information is imperfect information. However, the information improves as sampling size increases. It is possible for "less to be more." Behavioral information can be more useful than strategic information, especially if poor strategy gives rise to successful behavior (e.g. by luck of environmental circumstance). The moral: don't conflate genotype and phenotype. Conflation of genotype and phenotype in evolutionary game theoretic models leads to dramatically altered system behavior.

VII. IMPLICATIONS AND FUTURE WORK

As we have demonstrated, the genotype-phenotype conflation in evolutionary game theoretic models is problematic. It clearly matters whether agents imitate the *strategy or* the *behavior* of other more successful neighbors. The imitation of *strategy* can be thought of as a *perfect* information scenario. The imitation of *behavior*, on the other hand, can be thought of as an *imperfect* information scenario. The difference in results becomes smaller as sample size increases. Neighborhood size is one way to vary sample size. If your neighborhood is very small, you have a very small sample from which to

---

[10] Without mutations, this number would be 0.0, whereas with mutations, it would hover slightly above 0.0.

[11] We make no attempt here to offer an analytic solution because it would probably prove rather tedious. On the undecidability in the spatialized prisoner's dilemma, see Grim (1997).

learn or evolve; you are limited to very *local* information. If your neighborhood is the entire population, on the other hand, then your sample is *global*.

An alternative approach to increasing sample size in imperfect imitation scenarios is to increase how far into the past an agent remembers its interactions with another agent. Consider coin flips. The more times you flip an unbiased coin, the closer the observed frequency of heads will approach the actual probability (genotype) - one half of the total number of coin tosses. Similarly, with our imitation of behavior model (model B), information improves as an agent observes more instances of neighbor behavior.

The results presented in this paper may also have implications for political philosophy, in particular for the concept of desert in the economic market. Political philosophers argue about whether players in the economic market deserve their market returns. The most prominent argument for economic desert is that successful players provide social utility by fulfilling needs. The concern, however, is that there is too much luck in the market place.[12] Along these lines, our work suggests that political philosophers might be conflating good strategy with a good phenotype. A good phenotype (economic success) does not necessarily entail a good genotype (underlying strategy), a point usually overlooked by those who argue that the notion of merit belongs in the sphere of the market. Of course saying this does not mean that there is no economic entitlement. Entitlement and desert are distinct concepts, as is evidenced by the lottery, and our work says nothing about economic entitlement. For example, you are entitled to your lottery winnings even if you didn't do anything to deserve them. Similarly, you might be entitled to your market returns even if you do not necessarily always deserve them.

The treatment of genotype and phenotype as two mutually exclusive influences on strategy and behavior was artificially constructed to illustrate a point. In more realistic situations, observation of behavior probably influences, but does not replace, an agent's pre-existing internal strategy. In this case, a learning rule that captures incremental change in strategy upon observation of other successful agents' behavior, may improve model realism. We did pilot a primitive learning rule in this model and developed a preliminary framework for implementation of Bayesian updating. Our preliminary work indicated that overall system behavior is highly sensitive to the effectiveness of the learning rule. Given a perfect learning rule, it would not matter if agents are imitating behavior or strategy.[13] However, a thoughtful examination of agent learning and overall system behavior is a salient direction for future work.

Environmental information is another interesting and important consideration. One agent's environment might differ from that of another agent, resulting in different expressions of strategy/ genotype. A given agent, for example, might happen to have more cooperative neighbors. The optimal strategy, then, may be to imitate those more successful neighbors *with a similar environment* (more cooperative neighbors). Because our aim here was to simply illustrate the point that imitation of behavior and imitation of strategy can lead to very different results, our model does not implement this. It would, however, be interesting to see if implementing this sensitivity to the environment in our model would dampen the difference in overall system behavior.

---

[12] See Hayek for a philosophical argument along these lines.

[13] See the appendix for how a Bayesian agent would learn the cooperate/defect parameter of another agent.

VI. REFERENCES

Axelrod, R. (1984), *The Evolution of Cooperation*, New York: Basic Books.

Grim, P. (1995), “The Greater Generosity of the Spatialized Prisoner’s Dilemma.” Journal of Theoretical Biology 173, 353-359.

Grim, P., Mar, G., St. Denis, P. (1998), *The Philosophical Computer: Exploratory Essays in Philosophical Computer Modeling*, MIT Press/Bradford Books, §4.4

Grim, P. (1997), "The Undecidability of the Spatialized Prisoner's Dilemma," *Theory and Decision* 42, 53-80

Hayek, F. (1960), *Constitution of Liberty*, Routledge, chapters 5 and 6

Lewontin, R., "The Genotype/Phenotype Distinction", *The Stanford Encyclopedia of Philosophy (Spring 2007 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2007/entries/genotype-phenotype/>.

Nowak, M. A. and May, R. M. (1992) "Evolutionary Games and Spatial Chaos," *Nature*, **359**(6398), 29 October, pp. 826-829.

Nowak, M. A. and May, R. M. (1993) "The Spatial Dilemmas of Evolution," *International Journal of Bifurcation and Chaos*, **3**, pp. 35-78.

Nowak, M. A. and Sigmund, K. (1992) "Tit For Tat in Heterogenous Populations," *Nature*, **359**, pp. 250-253.

Wilensky, U. (1999), NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL.

Appendix A:
Bayesian learning of cooperation/ defection strategies

The strategy of a given player is given by the probability of cooperation π. The behavior of the player is described by the Bernoulli distribution

$$\begin{aligned}\mathbb{P}(\text{cooperate}) &= \pi\\ \mathbb{P}(\text{defect}) &= 1-\pi.\end{aligned}$$

A Bayesian would put a prior on π. The simplest prior to use is the β-distribution with parameters α and β:

$$p(\pi) = \text{beta}(\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\,\pi^{\alpha-1}\,(1-\pi)^{\beta-1}.$$

By Bayes' rule if we observe that the agent cooperates, the posterior becomes

$$p\,(\,\pi\,|\,\text{cooperate}\,) = \frac{p\,(\,\text{cooperate}\,|\,\pi\,)\;p(\pi)}{p(\text{cooperate})} \propto \pi\;\pi^{\alpha-1}(1-\pi)^{\beta-1}$$

from which we infer that the posterior is also a β-distribution, but with paramaters $\alpha+1$ and β:

$$p\,(\,\pi\,|\,\text{cooperate}\,) = \text{beta}(\alpha+1,\beta) = \frac{\Gamma(\alpha+\beta+1)}{\Gamma(\alpha+1)\Gamma(\beta)}\,\pi^{\alpha+1-1}\,(1-\pi)^{\beta-1}.$$

Similarly, observing that the agent defects will increase β by one:

$$p\,(\,\pi\,|\,\text{defect}\,) = \text{beta}(\alpha,\beta+1) = \frac{\Gamma(\alpha+\beta+1)}{\Gamma(\alpha)\Gamma(\beta+1)}\,\pi^{\alpha-1}\,(1-\pi)^{\beta+1-1}.$$

In this sense, the hyperparameters α and β can be regarded as "pseudocounts". It makes sense to set $\alpha=\beta$ to a small integer, for example, $\alpha=\beta=1$.