

Evolving to Generalize: Trading Precision for Speed

Cailin O'Connor

ABSTRACT

Biologists and philosophers of biology have argued that learning rules that do not lead organisms to play evolutionarily stable strategies (ESSes) in games will not be stable and thus will not be evolutionarily successful (Harley [1981]; Maynard-Smith [1982]). This claim, however, stands at odds with the fact that learning generalization—a behaviour that cannot lead to ESSes when modelled in games—is observed throughout the animal kingdom (Mednick and Freedman [1960]). In this article, I use learning generalization to illustrate how previous analyses of the evolution of learning have gone wrong. It has been widely argued that the function of learning generalization is to allow for swift learning about novel stimuli. I show that in evolutionary game theoretic models, learning generalization—despite leading to sub-optimal behaviour—can indeed speed learning. I further observe that previous analyses of the evolution of learning ignored the short-term success of learning rules. If one drops this assumption, I argue, it can be shown that learning generalization will be expected to evolve in these models. I also use this analysis to show how ESS methodology can be misleading, and to reject previous justifications about ESS play derived from analyses of learning.

- 1 *Introduction*
 - 2 *The Evolution of Learning*
 - 3 *The Approximation Game*
 - 4 *Learning Rules*
 - 4.1 *Herrnstein reinforcement learning and generalized reinforcement learning*
 - 4.2 *Long-term success*
 - 5 *Short-Term Success and Simulation*
 - 6 *Evolving to Generalize*
 - 7 *Conclusion*
-

1 Introduction

Stimulus generalization, or learning generalization, is a learning behaviour wherein an actor conditioned to one stimulus responds in the same way to perceptually similar stimuli.¹ This type of learning is extremely well documented.² It occurs across a wide variety of test subjects—mammals, birds, reptiles, amphibians, insects—across contexts, and across sensory modalities (Mednick and Freedman [1960]; Ghirlanda and Enquist [2003]). In evolutionary game theoretic models, however, learning generalization does not lead to the play of what are called ‘evolutionarily stable strategies’ (ESSes). One point that theorists have generally agreed on is that learning rules that do not lead organisms to play ESSes in games will not be stable and thus not evolutionarily successful (Harley [1981]; Maynard-Smith [1982]). Why this incongruity?

In this article, I will use the case of learning generalization to investigate how previous analyses of the evolution of learning have gone wrong. I point out that such analyses have largely ignored the short-term behaviour of learning rules. Learning generalization is standardly thought to be adaptive because it allows actors to quickly learn to respond to novel stimuli (Ghirlanda and Enquist [2003]). In other words, it is especially useful in the short term. I present evolutionary game theoretic models of learning generalization and show that, indeed, generalizing can be beneficial in these models in that it helps speed learning. Furthermore, if one considers evolutionary models of learning where the short-term behaviour of learning rules is important, it becomes clear that generalization can evolve. This supports the argument that previous analyses ignoring short-term learning were misguided. These results further inform game theory. Previous theorists used analyses of learning to argue that ESS behaviour should be seen in the real world. The work presented here indicates that such claims are overly hasty. Furthermore, this analysis lends credence to the idea that ESS methodology is often misleading.

The article will proceed as follows: In Section 2, I will discuss previous work on the evolution of learning. In Section 3, I will outline the ‘approximation game’, which appropriately models the class of scenario in which generalization is seen. In Section 4, I describe several learning rules where actors generalize to varying degrees. I go on to show that in the long run, rules that do not generalize outperform those that do in the approximation game. In Section 5, I present simulation results showing that despite the long-term success of non-generalized learning, under certain parameter settings higher levels of

¹ This behaviour was documented in the famous ‘Little Albert’ experiment. Watson and Rayner ([1920]) conditioned a nine-month infant to fear a white rat by frightening the child with loud noises whenever he touched the animal. The child subsequently showed similar fear reactions to a number of fuzzy stimuli, including a rabbit and a fur coat.

² Thankfully not with regard to infant fear response.

generalization can do significantly better in the short term. In Section 6, I show that in evolutionary game theoretic models where short-term learning is important, learning generalization can evolve. I conclude by discussing how this analysis informs game theory and evolutionary game theory.

2 The Evolution of Learning

Harley ([1981]) and Maynard-Smith ([1982]) use evolutionary game theoretic models to show that only certain sorts of learning rules should be expected to evolve. Without going into too much detail, these authors argue that only learning rules that lead to play of ESSes in games should be expected to persist in an evolutionary setting. An ESS is a strategy in a game that is robust against invasion by other strategies because it garners high payoffs for those using it.³ The arguments Maynard-Smith and Harley give are intuitively straightforward. Suppose that some learning rule does not lead to play of ESSes in games. A rule that does lead to play of an ESS will provide a higher payoff for those employing it. Then a learning rule leading to ESSes will be more evolutionarily successful than one that does not, and will be able to invade a population of those using a non-ESS learning rule.

This argument leads to a puzzle, however. Generalized learning cannot lead to play of ESSes in games (as I will show in Section 4). How does the observed ubiquity of learning generalization in the natural world square with these results?

The work of Maynard-Smith and Harley, of course, is not the end of the discussion of the evolution of learning. It has been pointed out by Smead ([2012]) that learning rules that take populations to ESSes have no advantage over static behavioural rules where the actor simply adopts ESS play rather than bothering to learn it.⁴ Furthermore, most models of the evolution of learning assume that learning will bear a greater cost than non-learning strategies (for cognitive architecture, time required to learn, and so on). This means that non-learning strategies that adopt ESS play will actually receive higher payoffs than rules that learn such play and so should be able to invade these learning rules. This point seems to create a worry about learning generally. If learning rules that do not lead to ESSes are unstable, and static behavioural rules can invade learning rules that do lead to ESSes, there are no stable learning rules at all (never mind ones that generalize).

³ To be specific, an evolutionary stable strategy, x_i , is one such that if $u(x_i, x_j)$ is the payoff of strategy x_i played against x_j : 1) $u(x_i, x_i) > u(x_i, x_j)$ or 2) $u(x_i, x_i) = u(x_i, x_j)$, and $u(x_i, x_j) > u(x_j, x_j)$ for all $x_j \neq x_i$.

⁴ Maynard-Smith ([1982]) was aware of this. Smead and Zollman ([unpublished]) find something similar. Smead ([2015]) also argues that learning rules that lead to equilibria in many cases should not be expected to evolve.

The usual response by biologists and philosophers of biology to worries of this sort is to argue that learning rules are primarily useful in situations where the environment exhibits some level of variability.⁵ In such environments, the argument goes, non-learning strategies get poor payoffs because the actors cannot respond to changing payoff structures by changing action. Actors that play an ESS in one situation, but cannot deal with changes to the environment, now do poorly against learners that reach this same ESS in the original situation and can re-adapt when necessary.

Something is amiss here, though. The arguments forwarded by Maynard-Smith and Harley explicitly depend on the following assumption: when modelling the evolution of learning one can ignore what happens in the short term. In other words, when associating fitnesses with learning rules, these authors do not consider payoff while the actors are learning. Instead, they look only at the payoffs of the long-term, stable strategies developed by learners. To date, most game theoretic work on the evolution of learning has shared this assumption.⁶

But if learning is most effective in a variable environment, to the extent that it should not be expected to evolve otherwise, this assumption is suspect. In a variable environment, an actor will be changing strategies and so may spend a significant amount of time playing strategies that are not stable, long-term outcomes of the learning process. If so, short-term behaviour should be important to the evolution of learning.⁷ In particular, if payoff in the short term matters, there should be selection pressure for learning rules that work quickly.

Biologists and psychologists have argued that the function of learning generalization is to allow organisms to quickly learn to respond to novel scenarios (Ghirlanda and Enquist [2003]). Furthermore, as mentioned, it should not evolve according to Maynard-Smith and Harley. As such, this learning behaviour is an excellent case to explore whether the intuitive argument I just gave—that short-term learning matters in an evolutionary context—is correct. In the rest of the article, I will present evolutionary game theoretic models of learning generalization. As I will show, when the short-term behaviour of learners is incorporated into evolutionary models, generalization will evolve for just the reasons that biologists and psychologists suggest. If short-term behaviour is ignored, on the other hand, generalization will not evolve. These results indicate that the intuitive argument is right, and that ignoring short term behaviour of learning rules can lead evolutionary analyses significantly astray.

⁵ See, for example, (Plotkin and Odling-Smee [1979]; Johnston [1982]; Maynard-Smith [1982]; Stephens [1991]; Godfrey-Smith [2002]; Dunlap and Stephens [2009]; Shettleworth [2009]).

⁶ There are some exceptions. Zollman and Smead ([2010]), for example, use interim strategies developed by learning rules to determine the fitnesses of actors in an evolutionary model.

⁷ Smead ([2012]) points out something similar. Empirical observations about, for example, death rates in young birds also confirm the important of learning speed in animals (Shettleworth [2009]).

3 The Approximation Game

Learning generalization occurs when an organism applies behaviour that was successful in one scenario to a perceptually similar scenario. What this means is that an appropriate model to explore the evolution of this phenomenon will need to include similar scenarios for the actor to potentially generalize over. In order to do this, I introduce the approximation game.⁸ The approximation game involves one actor and occurs in two stages. In the first stage, a state of the world is chosen probabilistically by nature or some exogenous force. In the second stage, the actor observes this state of nature and chooses an act. The state/act combination then determines what sort of payoff the actor receives. In order to model the type of scenario in which generalization evolves, the possible states of the world are assumed to bear similarity relationships to one another. This is done by treating these states as existing in a metric space where distance represents similarity. For example, an approximation game might have three states (1, 2, and 3) existing on a line. If state 1 is closer to state 2 than to state 3, it is assumed that state 1 is more similar to state 2.⁹

For each state of the world in the approximation game, it is assumed that there is some ideal act that, should the actor choose it, will give a perfect payoff.¹⁰ It is also assumed that acts will receive similar payoffs in similar states. In the previous example, in state 1 the actor would achieve a perfect payoff by choosing act 1. But she would also obtain a good payoff for choosing act 2. Her payoff for choosing act 3 would be less good. One simple way to model this is to determine payoff using a function that takes as input the distance between the state and the act.¹¹ For the purposes of this article, unless otherwise specified it will be assumed that the actor's payoffs are strictly decreasing with distance between state and act.

Figure 1 shows the simplest approximation game of interest—the one described above. The central node of the figure represents the starting point of the game, where nature chooses a state (S1, S2, or S3). The probabilities that each state is chosen by nature are fixed at p , q , and $1 - p - q$. The three decision nodes, labelled 'A' for actor, represent the possible choices of act in each state (A1, A2, or A3). Payoffs for each state/act combination are shown

⁸ This model should more properly be called the 'approximation problem' because it is a one-player decision problem rather than a multi-player game. Decision problems, however, are formally identical to one-player games. For this reason, the relevant results on the evolution of games directly bear on decision problems, and results from the problems investigated here can be used to inform evolutionary game theory. For simplicity sake, then, I use the language of game theory, and not decision theory, to describe the model used.

⁹ Note that this is similar to the sim-max game, introduced by Jäger ([2007]) to model signalling in situations where states of the world bear similarity relations to one another.

¹⁰ For simplicity sake, acts will always be labelled by the state they are most appropriate for, that is, act 1 will be the ideal act for state 1 and so forth.

¹¹ This is a useful way to understand payoff in these games. It is more precise to say that a payoff is defined for each state-act pair, and this payoff is chosen using such a function.




Figure 1. A 3-state/3-act approximation game with payoffs 1, δ , and ϵ for distance of 0, 1, and 2 between state and act. The game begins with the central node labelled ‘N’ for nature and continues to the three decision nodes labelled ‘A’ for actor.




Figure 2. Two examples of state spaces for an approximation game. Diagram (a) shows a game with four states modelled on a line. Diagram (b) shows a game with eight states modelled in a plane.

at the final nodes. It is assumed that $0 < \epsilon < \delta < 1$ (payoff decreases strictly in distance between state and act, but is always positive). It is also assumed that p and q are strictly positive and $p + q < 1$ (that every state is played with positive probability).

Figure 2 shows some possible state spaces for approximation games. Diagram (a) represents the state space of a game like the one just outlined, that is, modelled on a line, but with four states. Diagram (b) shows a game with a two-dimensional state space.¹² Approximation games with state spaces of any dimensionality are possible, though this article will only consider the simplest ones—those where states are modelled on a line. For the purposes of this article, these spaces are best understood as representing perceptual

¹² Note that games with state spaces of higher dimensionality can be used to model cases where an actor is responding to states with multiple properties varying along different dimensions.

similarity spaces.¹³ In other words, the states of the world in the game correspond to perceptual states. This is a useful interpretation of the model as learning generalization happens over perceptually similar states. It also avoids sticky issues around how or whether external states are similar to each other.

Most of the approximation games considered in this article will have a few properties that bear mentioning. First, they will have considerably larger state spaces than the game described above. The reason for this is that in real world learning scenarios, the number of possible states of the world is often extremely large. This is certainly true under the interpretation of the game here—that the actor is responding to perceptual states. Consider, for example, the number of discriminable colours picked out by the human visual system, or the number of distinguishable smells. Furthermore, as I shall show later in the article, considering games with large state spaces is relevant for understanding why generalized learning might evolve. Second, in the games considered, payoff loss over distance will usually be modelled with a Gaussian function. This function is used because it is always positive and strictly decreasing. These attributes make it particularly tractable from a modelling perspective. While this choice may seem arbitrary, the analytic results presented are robust under choice of function as long as it is strictly decreasing.¹⁴ I will call the Gaussian just described the ‘payoff Gaussian’ as it determines the degree to which an approximate match of state and act will lead to payoff for the actor.

As noted, for every state of an approximation game, there is one ideal act. A strategy for a game defines an act in every possible state.¹⁵ What this means is that there is a single, optimal strategy for every approximation game in which the actor always picks the correct act for the state. The existence of a single optimal strategy is significant from an evolutionary standpoint. Under the replicator dynamics, the most common model of evolutionary change in evolutionary game theory, a population playing the approximation game will evolve to take this strategy in every case. For this reason, the approximation game would not usually be of much interest to evolutionary game theorists—it is immediately obvious what behaviour will be adopted by a population evolving to play it. However, as I will argue in the next section, an organism learning to respond to this game, and employing generalizing learning, will not develop the optimal strategy.

¹³ See (Gärdenfors [2000]) for more on such spaces. See (Krantz *et al.* [1971]) for how such spaces can be built using experimental data.

¹⁴ O'Connor ([2014a]) also found that results in simulations of related signalling games were robust under choice of function for payoff loss modelled as linear, quadratic, or decreasing in steps.

¹⁵ Again, while the term that technically should be used here is ‘choice’, because this is a one-player problem, I use ‘strategy’ to avoid confusion. Once again, nothing hangs on this distinction.

4 Learning Rules

4.1 Herrnstein reinforcement learning and generalized reinforcement learning

In evolutionary game theory, learning dynamics, unlike evolutionary dynamics, are taken to model the emergence of learned individual behaviours over the course of an organism's lifetime, rather than the emergence of evolved population behaviours over the course of evolutionary time. Herrnstein reinforcement learning, first proposed by Roth and Erev ([1995]), is so named in reference to R. J. Herrnstein's psychological work on learning, which motivates the model (Herrnstein [1970]).¹⁶ This learning rule has been widely used in evolutionary game theory because (1) it is psychologically natural, that is, based on observed learning behaviour, and (2) it makes minimal assumptions about the cognitive abilities of the actors. This means that behaviours that emerge under this rule can be assumed to be available to cognitively simple animals. In this case, because generalized learning is seen in a wide variety of animals, including those with minimal cognitive abilities (Mednick and Freedman [1960]), Herrnstein learning is an appropriate starting place to model it.

The basic assumption that underlies reinforcement learning rules is that actors will be more likely to repeat successful behaviour. In other words, they reinforce this behaviour. In a simulation of these rules, actors engage in a game many times, at each step reinforcing successful behaviour and thus improving their strategies. Herrnstein learning can be described using the following analogy: In the context of the approximation game, imagine that for each state of the world, the actor has an urn into which is placed one coloured ball for each possible act available. In the first round of learning, nature selects a state of the world and the actor draws a ball from the urn for that state. The colour of the ball determines which act the actor will take. If the act is successful, the actor returns the drawn ball to her urn and then reinforces their tendency to take that act in that state by adding a ball (or two, or half a ball, and so on) of the same colour to that urn. The reinforcement is proportional to the success of the act, that is, the higher the success the greater the reinforcement. For our purposes, the amount of reinforcement will always be equal to the payoff achieved by the actor in each step of the simulation. At the beginning of a simulation using Herrnstein learning, an actor uses all her acts with equal probability, as they have one of each type of ball in each urn. As play progresses and successful acts are reinforced, the actor becomes increasingly likely to choose these acts. In the limit, the actor's strategy

¹⁶ This learning rule is also sometimes called 'Roth-Erev' or 'Vanilla' reinforcement learning.




Figure 3. A representation of how the payoff and reinforcement Gaussians determine reinforcement in an approximation game evolved using a GRL rule.

may, under the right circumstances, converge to a successful one. In other words, the actor will use this strategy with probability approaching one.¹⁷

Generalized reinforcement learning (GRL) builds on the Herrnstein reinforcement learning model.¹⁸ Under GRL rules, successful acts are reinforced, but they are also generalized, that is, reinforced for other, similar states of the world. In other words, and to continue the urn analogy, when an actor draws a coloured ball from her urn for a state and takes a successful act, she adds balls of the same colour to that urn, but also adds balls of that colour to the urns for similar states. It must be specified, for these rules, the degree to which generalization occurs. How many other states are reinforced? How much reinforcement occurs in those states? For the purposes of this article, generalization will be determined using a Gaussian function. To be clear, a model of an approximation game evolved using GRL employs two Gaussian functions. The payoff Gaussian, introduced above, determines the level of payoff based on how accurate the act chosen is for the state. The second Gaussian determines to what degree this payoff is generalized taking as input the distance between the state of the world and the state to be reinforced. I will call this second Gaussian the ‘reinforcement Gaussian’.¹⁹ **Figure 3**

¹⁷ For more on this and other learning dynamics see (Huttenberger and Zollman [2011]). For extensive work on Herrnstein reinforcement learning and variations of it in signalling games (which are in some ways similar to the approximation game), see recent work by Barrett ([2007], [2009]); Barrett and Zollman ([2008]).

¹⁸ This learning rule was first outlined by O’Connor ([2014a]). Roth and Erev ([1995]) look at a learning rule that incorporates a slight amount of generalization in a similar way to GRL. They interpret this aspect of the learning rules as persistent error.

¹⁹ Ghirlanda and Enquist ([2003]) argue that generalization is best modelled in many cases by a Gaussian function, suggesting that the choice of a Gaussian as the reinforcement function here is a natural one. Furthermore, Shepard ([1987]) argues that the specifics of how an actor learns to generalize may not be particularly important in determining subsequent behaviour.

represents the way these two functions determine reinforcement in an approximation game evolved using a GRL rule.

A model of an approximation game evolved using these learning rules will have five relevant parameters. The first is the size of the state space of the game. The second and third are the height and standard deviation of the payoff Gaussian. These control the level of payoff for perfect coordination in the approximation game (the height) and the degree to which an actor receives payoff for imperfect action in the game (the standard deviation). The fourth parameter is the standard deviation of the reinforcement Gaussian.²⁰ Variations of this parameter correspond to GRL rules with different degrees of generalization. In models where Herrnstein reinforcement learning is used, this parameter will not apply. It can be noted, though, that Herrnstein learning is a limiting case of GRL as the width of the reinforcement Gaussian approaches zero. The fifth relevant parameter will be the length of trial for simulations of these models. This parameter will control the number of times the actor plays the approximation game and updates her strategies.

4.2 Long-term success

One way to explore the evolution of generalized learning is to compare learning rules with different levels of generalization, like GRL and Herrnstein reinforcement learning, to see if high levels of generalization can outperform lower levels in these models. One method for doing this is to consider convergence outcomes of the models just described. When this is done, however, it becomes clear that in the long term, Herrnstein reinforcement learning can always outperform GRL in the approximation game.

Laslier *et al.* ([2001]) show that a single actor employing Herrnstein learning in a stationary environment, that is, where payoffs remain constant, in the long run will always learn to play the act that receives the highest expected payoff.²¹ This result can be applied to each state in the approximation game. To do so requires that each state be a stationary environment, which is the case given that the payoffs in the approximation game do not change. It also requires that each state be selected infinitely often as the length of learning goes to infinity, which is also the case as each state in the approximation game has a strictly positive probability. Thus these results indicate that in the long run, for each state in the approximation game, the act of an agent employing Herrnstein reinforcement learning will converge to the optimal one. For the entire game, then, the strategy of the actor will converge to the optimal strategy. In the long run, the actor will take the perfect act in every state in the

²⁰ The height of the reinforcement Gaussian is determined by the level of payoff.

²¹ In other words, as the learning time goes to infinity, the probability with which the actor chooses non-optimal acts goes to 0.




Figure 4. A 2 state/2 act approximation game with payoffs 2 and 1 for distance of 0 and 1 between state and act. The game begins with the central node labelled ‘N’ for nature and continues to the two decision nodes labelled ‘A’ for actor.

approximation game if using Herrnstein learning. This result holds for an approximation game of any finite size.

What happens to the strategy of an actor using a GRL rule in the approximation game in the long run? Unlike Herrnstein learning, GRL rules will not converge to the optimal strategy and, in fact, the level of generalization will determine a bound of accuracy that a player will not be able to surpass. This bound of accuracy will in turn determine a bound on the payoff success an actor can achieve. The intuitive reason for this is that if an actor were able to converge to the perfect act in one state, she would simultaneously prevent convergence in neighbouring states by generalizing the same act to them.

One can show this by solving for the consistent, limiting probabilities of acts for a model of the approximation game evolved using a GRL rule. This is done by finding the distribution of reinforcements in a game where the probability of an act being selected in one round of simulation is equal to the probability of it being selected in the next round. Consider a toy model of the approximation game with two states and two acts. Suppose that in each state the payoff for the perfect act is 2 and for the other act is 1. Assume that states of the world are equiprobable.²² This game is pictured in Figure 4, which should be read like Figure 1. Also consider a simple form of GRL where successful acts are reinforced in the state of the world by the amount of the payoff and in the other state by that amount multiplied by α where $0 \leq \alpha \leq 1$. In this simple model, α determines the level of generalization. A high α means that success will lead to strong generalization in the other state of the world; a low α will mean that generalization is weak. If α is equal to 0.1, the consistent, limiting probabilities of this game are such that the actor selects the more successful act

²² This degenerate approximation game is not generally an interesting one as it is formally the same as a game with no similarity structure over the payoffs. It is useful, however, as a simple case to consider GRL.




Figure 5. Success levels for a 2 state/2 act approximation game with various levels of generalization (α). The y -axis tracks success and the x -axis represents length of the trial where each value x is 10^x runs.

in each state with probability $5/6$ and the other act with probability $1/6$. It is possible (though increasingly difficult) to calculate such limiting probabilities for larger games and more complex generalization rules.

One can further explore this phenomenon through simulation. It is easy to show what happens in this toy model at the two bounds of α . If one sets $\alpha = 0$, the learning rule is the same as Herrnstein learning and so converges to perfect behaviour. If one sets $\alpha = 1$, the actor fully generalizes. In other words, if she reinforces act 1 in state 1 by 0.43, she will also reinforce act 1 in state 2 by 0.43, and so on. This complete generalization of success means that reinforcement levels for the actors will always be identical in the two states of the world. Because actors will not be able to learn to condition their acts on which state has been selected, every attainable strategy (those where the probability for each act is the same in both states) will get an expected payoff of 1.5, the same as choosing by chance.

For intermediate levels of α , simulations of the toy model show that the actor eventually reaches a level of accuracy, and thus success, that is bounded by the level of generalization. The lower the generalization, the greater the success. In Figure 5, success rates are shown for a simulation of this game for α ranging from 0 to 0.3 and α equal to 1. In each case, success is calculated by dividing the expected payoff for the actor given her learned strategy by the perfect possible expected payoff (which, in this case, is two).

$$\text{Success} = \frac{\text{expected payoff given learned strategy}}{\text{perfect possible expected payoff}}.$$

Each line represents the success rate of a simulation over time for a different level of generalization. Darker lines represent lower levels of generalization. Rates were averaged over fifty runs of simulation. As should be clear from Figure 5, for each level of generalization, the success of the simulation reaches some upper bound and stays there. Note that time is presented logarithmically. The reason for this bound on success has already been laid out. When the actor generalizes, success in one state means that an act will be taken with greater probability in other states where it is less successful.

The results from these toy models can be extended to larger approximation games, since in every larger game, reinforcement in neighbouring states will prevent convergence in the same way as it does in a two state model.²³ Thus these results indicate that in the approximation game, over the long run, low levels of generalization will outperform high levels of generalization from a payoff perspective, and in particular Herrnstein learning will outperform any GRL rule. The single optimal strategy provides the highest possible level of payoff in the game, and so learning to use any other strategy will be strictly worse. Importantly, the optimal strategy in an approximation game is always the unique ESS. Therefore, GRL is unable to learn ESSes in this game, while Herrnstein is guaranteed to do so.

Furthermore, although this analysis only addresses approximation games, it may be extended to some other games, including ones with multiple players. O'Connor ([2014a]) obtained similar simulation results in sim-max games, which are a variation on the Lewis signalling game where the state space has the same similarity structure as the approximation game. Unlike approximation games and sim-max games, most games do not have several possible states and so it is not possible to evolve them using GRL. For those games that do, though, if an actor generalizes over states she will only be able to achieve optimal behaviour if the acts so generalized are ideal for all the states they are generalized to. Otherwise, generalized learning will lead to reinforcement of sub-ideal acts and thus to sub-optimal behaviour, preventing play of ESSes.

5 Short-Term Success and Simulation

As I will outline in this section, there is a tension that can arise between the two *desiderata* a learning rule should meet—working quickly and developing behaviour that obtains the highest possible payoff.²⁴ While low generalization

²³ To see why this is the case, consider two states of any larger approximation game. Use the reinforcement Gaussian for this larger game to define α as above (the proportion of reinforcement on a neighbouring state). It has been shown that this smaller system cannot reach an optimal strategy and so the larger system it is a part of cannot either.

²⁴ This has been widely observed in other fields. It has been argued in psychology that ‘fast and frugal’ decision heuristics, which allow actors to make decent strategies quickly and easily, are adaptive, despite the possibility that they lead to irrational or sub-optimal behaviour

learning outperforms high generalization learning eventually, the very property that prevents high generalization rules from approaching optimal behaviour is the one that allows them to outperform low generalization rules in the short term. I will illustrate this argument using simulation results showing that in trials of the approximation game, high levels of generalization can outperform low levels under certain parameter settings. In particular, high generalization does best when states of the world are numerous, when trials are short, and when the payoff Gaussian (modelling how accurate an actor must be to get a good payoff) is wide. This result confirms intuitive arguments about the benefits of learning generalizations.

All the results presented in this section were generated using models where payoff and reinforcement were calculated with Gaussian functions. Each trial of a parameter setting was run fifty times and reported results are averages of these. The parameters that varied were the size of the state space, the length of the trial, the standard deviation of the reinforcement Gaussian, and the standard deviation of the payoff Gaussian.²⁵ The state spaces considered were of size 100, 200, 300, 400, and 500. The lengths of trial were 1000, 10,000, 100,000, and 1 million runs. The reinforcement Gaussian standard deviations were 5, 10, 15, 20, and none (Herrnstein learning). And lastly, the payoff Gaussian standard deviations were 1, 5, 10, 15, and 20.

Figure 6 shows the success rates (calculated as they were in the previous section) for one set of these trials—those where the payoff Gaussian had a standard deviation of 10. The x -axis of the figure represents the length of trial (ranging from 1000 runs to 1 million). The z -axis tracks the size of the state space (from 100 to 500), and the y -axis tracks average success of the trials. Each surface shown represents results for one reinforcement width parameter setting. In other words, each surface corresponds to one learning rule and these rules vary with respect to generalization. The black surface represents the highest levels of generalization (a reinforcement Gaussian with a standard deviation of 20) and successively lighter surfaces represent lower and lower levels of generalization.

As is evident in the figure, each level of generalization considered outperforms the others for some region of parameter space. The rule with the highest level of generalization (black) outperforms the others in the area of parameter

(Gigerenzer and Selten [2001]; Gigerenzer and Gaissmaier [2011]). Generalized learning can be thought of as a learning rule that leads to making decent, if sometimes inaccurate, decisions quickly. In machine learning, much work has been done on learning models that generalize from limited input to make predictions in novel scenarios. Similar trade-offs between speed and accuracy are found in these models (Hastie *et al.* [2005]).

²⁵ Height of the payoff Gaussian was always 2.




Figure 6. Average success levels for various parameter settings for an approximation game with a payoff Gaussian of standard deviation 10 evolved using GRL and Herrnstein reinforcement learning. Results are averaged over fifty runs of each setting.

space where trials are short and the number of states of the world is large. Herrnstein learning (the lightest surface) performs best in the longest trials and when states of the world are fewer. These results should not be surprising. In a short trial with many states of the world, there is not enough time for the actor to learn ideal actions in each state, so a learning rule that allows success to be generalized does better. When an actor has a long time to learn, more precise strategies can be developed using low generalization rules and so these do better. Similar results were obtained for the other payoff Gaussian values with the slight difference that in games where approximate action was successful (wide payoff Gaussians), higher generalization could perform better. In extreme cases of games with very narrow payoff Gaussians, approximate actions do not receive a good payoff. Generalization thus does not help the actor in this case, because only precise strategies will be successful.

Real world learners do not use learning strategies that exactly mimic those used in the models here. In order to strengthen these results, I investigated their robustness across learning rules. Under reinforcement learning with punishment, actors reinforce successful acts for the state of the world (and for similar states under the generalized version), and simultaneously punish, or

decrease the reinforcement level, for that act in other states.²⁶ The results of simulations for these rules were highly similar to those presented in the last section. I also explored a learning rule outlined by Barrett ([2014]), which I call 'Barrett learning'. This rule is in some ways similar to adjustable reference point with truncation learning introduced by Bereby-Meyer and Erev ([1998]). Actors using this rule discount past experience compared to more recent experience. Results were, again, very similar to those presented in this section.

It should be noted that the results presented in this section are not particularly surprising given previous results from machine learning, and previous observations from psychology and biology about the benefits of generalization. As we will see in the next section, however, generating similar results in an evolutionary game theoretic model is useful in that it allows us to discuss the motivating problem presented in Section 2: why do previous evolutionary game theoretic analyses of learning predict that rules like GRL should be unable to evolve if generalization is so ubiquitous?

6 Evolving to Generalize

At this point it has been established that high generalization learning can perform well in the approximation game when time is limited and states are numerous, despite the fact that only non-generalized learning leads to optimal behaviour. How, it will now be asked, do these results inform the evolution of learning generalization?

The larger question at hand, remember, is whether or not it is problematic to assume that the short-term behaviour of learning rules does not matter in evolutionary analyses. In order to assess this using the case of learning generalization, let us consider an evolutionary model where the environment for the actor changes regularly, meaning that speed of learning may be evolutionarily relevant. The replicator dynamics are the most commonly used model of the evolutionary process in evolutionary game theory and will be employed here. These dynamics assume that actors using strategies that receive higher payoffs will replicate more successfully than actors using strategies that receive lower payoffs.²⁷ In populations modelled under these dynamics, high payoff strategies tend to proliferate. In the approximation game in particular, because there is only one player, the learning rule that will evolve under the replicator dynamics is simply the one that gets the best payoff.

²⁶ There is experimental evidence supporting the use of rules where actors punish or forget strategies, that is, they sometimes decrement their reinforcements. See (Bereby-Meyer and Erev [1998]), for example.

²⁷ The replicator equation determines how proportions of strategies in a population change under the replicator dynamics. This equation states that $\dot{x}_i = x_i(f_i(x) - \sum_{j=1}^n f_j(x)x_j)$ where x_i is the proportion of a population playing strategy i , $f_i(x)$ is the fitness of type i in the population state x and $\sum_{j=1}^n f_j(x)x_j$ is the average population fitness in this state.

Consider a model where a population of actors learns to play an approximation game using either Herrnstein learning or various GRL rules. One can think of the actors' strategies as now consisting in which learning rules to adopt. The payoffs associated with each learning rule will be the expected payoffs for the behavioural strategies that these various learning rules develop in simulation. Now suppose that at regular intervals, the population encounters a new approximation game (one where the actors encounter new states and must associate them with new actions). If these intervals of learning are short enough, under the replicator dynamics this population will evolve to use a GRL rule rather than Herrnstein learning. This is the case because, as shown in the previous section, generalizing rules will lead to higher payoffs for the actors over a short timescale. And, as pointed out, for an approximation game the replicator dynamics will always select whichever behaviour receives the best payoff. To give an example, suppose that actors in the population play approximation games with 100 states, and that they switch games every 1000 rounds. If the initial population contains the learning rules considered in the last section (Herrnstein learning and GRL with reinforcement Gaussians of widths 5, 10, 15, and 20), GRL with a reinforcement Gaussian of width 10 will evolve. In other words, when the environment varies, generalization can evolve.

One might worry that in the model just described actors begin their learning processes anew when the environment changes rather than having to forget currently developed actions. To alleviate this worry, I also considered models of populations in changing environments where actors must forget previously learned strategies when the world changes. I found that under a wide range of parameter settings, generalization evolved.²⁸

Furthermore, there is a feature of learning situations that I have not discussed yet that makes generalization relatively more important and more successful in real world scenarios with numerous states. In the approximation game, every possible state of the world has its own ideal act. In reality though, for highly similar states it will often be appropriate for an organism to take the same act, in which case generalization will be more effective than the models here predict (Shettleworth [2009]). To further elucidate this claim, consider a scenario where a bird is learning to interact with blackberries. Imagine a model of this scenario. The state space of this model would have hundreds (thousands?) of states varying along multiple dimensions of perceptual space—smell, size, colour, shape, and so on—but the birds would only have two available acts—eat and not-eat. Generalization, in this case, will only lead to sub-optimal behaviour for states right at the boundary between edible and

²⁸ These results are not presented here as the description of these models is lengthy and the results are unsurprising.

inedible berries. For all the other possible states, generalizing will be completely successful. In models of this scenario, Herrnstein learning will still lead to ESS play while GRL will not, but the benefits of Herrnstein learning are only relevant for a small proportion of states, while GRL provides more significant benefits for most of the state space. In other words, the window of time during which GRL is a more successful learning rule is longer, making it more problematic to ignore short-term learning behaviour in evaluating the evolution of generalization.

In the evolutionary models presented above, the learning rule that evolves strictly outperforms the other learning rules from a payoff perspective, and so satisfies the definition of an ESS (if one treats a choice of learning rule as a choice of strategy). It would be strange to say that GRL rules are evolutionarily stable, though. In principle, given this set-up, any learning rule (like GRL) that has not gotten to the optimal outcome in the short time period could be outperformed, and so invaded, by a learning rule that does better in that same time period.²⁹ This, however, does not really matter. The point is not that a particular rule for generalization will be stable, but rather that this type of stability analysis ignores some of the most evolutionarily relevant features of learning rules, in this case a need for speed. Maynard-Smith ([1982]) and Harley ([1981]) are not wrong in thinking that there should be selection pressure for rules that learn ESSes, just wrong in thinking that this is the only, or the most important type, of selection pressuring bearing on learners.

7 Conclusion

To conclude, I will discuss how the results of this article inform game theory and evolutionary game theory; but first, a word should be said about the proposed interpretation of the state spaces of approximation games.

I pointed out in Section 2 that these state spaces should be thought of as perceptual rather than external because generalization happens over perceptually similar states. Given that similarity is built into the approximation game

²⁹ In fact, the real world behaviour of learning discrimination points towards a possibility for such an improved rule. Previous investigations into animal learning indicate that when it is relevant from a payoff perspective for organisms to discriminate between states, they learn to do so (Mackintosh [1974]). In fact, generalization and discrimination can be seen as two sides of a coin: the former allows animals to extend successful behaviours to possibly relevant scenarios, the second allows animals to trim these behaviours back if they are not applicable (Shettleworth [2009]). This combination of behaviours could be modelled with a learning rule that combines the best aspects of GRL and Herrnstein reinforcement learning. Learners begin by generalizing, but eventually stop generalizing and develop more precise strategies. In fact, the models developed here help illuminate why learning discrimination is important: it helps organisms avoid sub-optimal behaviours developed when generalizing and can allow actors to move closer to ESSes.

through the payoff structure, however, this interpretation assumes that perceptually similar states will always get similar payoffs when responded to with similar actions. At first consideration, this assumption may seem problematic. It should be noted, though, that perceptual similarity structures themselves evolve. O'Connor ([2014b]) argues that in models of the evolution of perceptual categorization, real-world states that actors can respond to in similar ways evolve to be perceptually similar. If this is right, it may be reasonable to assume that perceptual similarity (usually) tracks payoff similarity. This line of thinking points to a way in which the exploration of generalization in this article is incomplete, though. Generalization happens over perceptual states, and will only be successful if the similarity structure of these perceptual states is arranged so that perceptually similar things can be reacted to similarly. In this way, the evolution of generalization arguably cannot be fully understood without also understanding the evolution of perceptual similarity.

I will now return to how this exploration of the evolution of learning generalization informs evolutionary game theory. First, and most importantly, the assumption that the short-term performance of learning rules can be ignored in evolutionary analyses is a bad one. This assumption is inconsistent with other assumptions made about the evolution of learning, in particular, that learning should be expected to evolve in variable environments. It is an assumption that matters because, as shown here, when the short-term success of learning rules is taken into account, evolutionary outcomes are significantly impacted. And, as this article shows, if this assumption is maintained, evolutionary game theoretic models are unable to account for the evolution of generalization. When the assumption is dropped, on the other hand, evolutionary game theoretic models can account for this highly successful real world learning behaviour. As such, this case illustrates how the long-term learning assumption is not just intuitively suspect, but can actually lead an evolutionary analysis significantly astray.

Past investigations into the evolution of learning rules have been used to justify assumptions about equilibrium play in game theory (see, for example, Maynard-Smith [1982]). The results here indicate that a better understanding of the evolution of learning does not support this justification. Although there should be selection pressure for learning rules to reach ESSes, there should also be selection pressure for rules that learn quickly. When these *desiderata* are at odds, as is the case with learning generalization, non-equilibrium behaviour should be expected in the real world. Even if real world actors eventually learn to discriminate between relevantly different states, and so mitigate the sub-optimal effects of generalization, while learning progresses (which should be a non-trivial proportion of the time if actors face heterogenous environments), non-equilibrium and thus non-optimal behaviour should be expected.

In recent years, the tradition of depending on ESS methodology in evolutionary analyses has come under fire. The results presented here are one more example of a case where a dynamical investigation reveals important insights into evolutionary processes that ESS analysis misses. As discussed, simply identifying which learning rules are evolutionarily stable in the sense that they lead to ESSes misses important differences between the processes that actors employing these rules undergo, and thus misses evolutionarily relevant information. This analysis thus gives further reason to be very careful when applying ESS methodology to complicated evolutionary scenarios.

Acknowledgements

Many thanks to Simon Huttegger, Michael McBride, Louis Narens, Kyle Stanford, Brian Skyrms, Elliott Wagner, and James Weatherall for comments on this work. Thanks to helpful audiences at ISHPSSB 2013, the Winter Q-Bio conference 2014, and the ABMP conference 2014, as well as at the Center for Philosophy of Science at the University of Pittsburgh. Special thanks to Rory Smead for his help at all stages of this project.

*Department of Logic and Philosophy of Science
University of California
Irvine, CA 92697, USA
cailino@uci.edu*

References

- Barrett, J. A. [2007]: 'Dynamic Partitioning and the Conventionality of Kinds', *Philosophy of Science*, **74**, pp. 527–46.
- Barrett, J. A. [2009]: 'The Evolution of Coding in Signaling Games', *Theory and Decision*, **67**, pp. 223–37.
- Barrett, J. A. [2014]: 'Description and the Problem of Priors', *Erkenntnis*, **79**, pp. 1343–53.
- Barrett, J. A. and Zollman, K. [2008]: 'The Role of Forgetting in the Evolution and Learning of Language', *Journal of Experimental and Theoretical Artificial Intelligence*, **21**, pp. 293–309.
- Bereby-Meyer, Y. and Erev, I. [1998]: 'On Learning to Become a Successful Loser: A Comparison of Alternative Abstractions of Learning Processes in the Loss Domain', *Journal of Mathematical Psychology*, **42**, pp. 266–86.
- Dunlap, A. S. and Stephens, D. W. [2009]: 'Components of Change in the Evolution of Learning and Unlearned Preference', *Proceedings of the Royal Society*, **276**, pp. 3201–8.
- Gärdenfors, P. [2000]: *Conceptual Spaces: On the Geometry of Space*, Cambridge, MA: MIT Press.

- Ghirlanda, S. and Enquist, M. [2003]: 'A Century of Generalization', *Animal Behaviour*, **66**, pp. 15–36.
- Gigerenzer, G. and Gaissmaier, W. [2011]: 'Heuristic Decision Making', *Annual Review of Psychology*, **62**, pp. 451–82.
- Gigerenzer, G. and Selten, R. [2001]: *Bounded Rationality: The Adaptive Toolbox*, Cambridge, MA: MIT Press.
- Godfrey-Smith, P. [2002]: 'Environmental Complexity and the Evolution of Cognition', in R. Sternberg and J. Kaufman (eds.), *The Evolution of Intelligence*, Mahwah: Lawrence Erlbaum, pp. 233–49.
- Harley, C. B. [1981]: 'Learning the Evolutionary Stable Strategy', *Journal of Theoretical Biology*, **89**, pp. 611–33.
- Hastie, T., Tibshirani, R. and Friedman, J. [2005]: 'The Elements of Statistical Learning: Data Mining, Inference, and Prediction', *The Mathematical Intelligencer*, **27**, pp. 83–5.
- Herrnstein, R. [1970]: 'On the Law of Effect', *Journal of the Experimental Analysis of Behavior*, **13**, pp. 243–66.
- Huttegger, S. M. and Zollman, K. J. S. [2011]: 'Signaling Games: Dynamics of Evolution and Learning', in *Language, Games, and Evolution*, Berlin Heidelberg: Springer-Verlag, pp. 160–76.
- Jäger, G. [2007]: 'The Evolution of Convex Categories', *Linguistics and Philosophy*, **30**, pp. 551–64.
- Johnston, T. D. [1982]: 'The Selective Costs and Benefits of Learning: An Evolutionary Analysis', in J. S. Rosenblatt (ed.), *Advances in the Study of Behavior*, Volume 12, New York: Academic Press.
- Krantz, D. H., Luce, R. D., Suppes, P. and Tversky, A. [1971]: *Foundations of Measurements, Volume 1: Additive and Polynomial Representations*, Mineola, NY: Dover Publications.
- Laslier, J. F., Topol, R. and Walliser, B. [2001]: 'A Behavioral Learning Process in Games', *Games and Economic Behavior*, **37**, pp. 340–66.
- Mackintosh, N. J. [1974]: *The Psychology of Animal Learning*, Oxford: Academic Press.
- Maynard-Smith, J. [1982]: *Evolution and the Theory of Games*, Cambridge: Cambridge University Press.
- Mednick, S. A. and Freedman, J. L. [1960]: 'Stimulus Generalization', *Psychological Bulletin*, **57**, pp. 169–200.
- O'Connor, C. [2014a]: 'The Evolution of Vagueness', *Erkenntnis*, **79**, pp. 707–27.
- O'Connor, C. [2014b]: 'Evolving Perceptual Categories', *Philosophy of Science*, **81**, pp. 840–51.
- Plotkin, H. C. and Odling-Smee, F. J. [1979]: 'Learning, Change, and Evolution: An Enquiry into the Teleonomy of Learning', in J. S. Rosenblatt (ed.), *Advances in the Study of Behavior*, Volume 10, New York: Academic Press.
- Roth, A. E. and Erev, I. [1995]: 'Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term', *Games and Economic Behavior*, **8**, pp. 164–212.
- Shepard, R. N. [1987]: 'Toward a Universal Law of Generalization for Psychological Space', *Science*, **237**, pp. 1317–23.

- Shettleworth, S. J. [2009]: *Cognition, Evolution, and Behavior*, Oxford: Oxford University Press.
- Smead, R. [2012]: 'Game Theoretic Equilibria and the Evolution of Learning', *Journal of Experimental and Theoretical Artificial Intelligence*, **24**, pp. 301–13.
- Smead, R. [2015]: 'The Role of Social Interaction in the Evolution of Learning', *British Journal for the Philosophy of Science*, **66**, pp. 161–80.
- Smead, R. and Zollman, K. [unpublished]: 'The Stability of Strategic Plasticity'.
- Stephens, D. W. [1991]: 'Change, Regularity, and Value in the Evolution of Animal Learning', *Behavioral Ecology*, **2**, pp. 77–89.
- Watson, J. and Rayner, R. [1920]: 'Conditioned Emotional Reactions', *Journal of Experimental Psychology*, **3**, pp. 1–14.
- Zollman, K. and Smead, R. [2010]: 'Plasticity and Language: An Example of the Baldwin Effect?', *Philosophical Studies*, **147**, pp. 7–21.