

Rule-Following and the Evolution of Basic Concepts

Jeffrey A. Barrett*†

This article concerns how rule-following behavior might evolve, how an old evolved rule might come to be appropriated to a new context, and how simple concepts might co-evolve with rule-following behavior. In particular, we consider how the transitive inferential rule-following behavior exhibited by pinyon and scrub jays might evolve in the context of a variety of the Skyrms-Lewis signaling game, then how such a rule might come to be appropriated to carry out inferences regarding stimuli different from those involved in the original evolution of the rule, and how this appropriation involves a step toward the evolution of basic ordinal concepts.

1. Introduction. There is a long philosophical tradition of puzzling over how one might learn and follow an abstract rule.¹ Part of the puzzle concerns how it is possible to apply a rule in a context that one never considered when initially learning the rule. Here we consider how rule-following behavior might evolve and how an old evolved rule might evolve to be appropriated to a new context. This will also provide an account of how basic concepts might coevolve with an increasingly general rule that characterizes the concepts.

q1

We start by considering the transitive rule-following behavior exhibited by pinyon and scrub jays and how such rule-following behavior might evolve in the context of a variety of the Skyrms-Lewis signaling game.² Then we consider how a rule evolved in this way might come to be appropriated to

*To contact the author, please write to: Department of Logic and Philosophy of Science, University of California, Irvine, CA 92697; e-mail: jabarret@uci.edu.

†I would like to thank Brian Skyrms and Cailin O'Connor for discussions on the topic of this article and an anonymous reviewer for thoughtful comments on an earlier draft.

1. See Kripke (1982) for a standard introduction to such puzzles.

2. See Lewis (1969) and Skyrms (2010) for basic descriptions of signaling games in the context of classical (Lewis) and evolutionary (Skyrms) games.

Philosophy of Science, 81 (December 2014) pp. 1–11. 0031-8248/2014/8105-00XX\$10.00
Copyright 2014 by the Philosophy of Science Association. All rights reserved.

a context different from that in which it initially evolved by simple reinforcement with punishment, an evolutionary process similar to the one that evolved the rule in the first place but much more efficient than evolving a new rule from scratch. This appropriation of the old rule to a new context involves the evolution of an analogy between the old and the new contexts. In this case, the evolution of such an analogy can be thought of as a step toward the evolution of basic ordinal concepts that fall under an increasingly general transitive ordering rule.³

2. The Curious Behavior of Jays. Pinyon jays *gymnorhinus cyanocephalus* and scrub jays *aphelocoma californica* are two of several species of animal that have evolved to exhibit transitive rule-following behaviors. Such rule-following behavior is illustrated in an experiment reported by Bond, Kamil, and Balda (2003).⁴

In their experiment, seven stimulus colors were arranged in a random linear order that was fixed for each bird. For one of the birds, for example, the order might be red, green, blue, magenta, yellow, cyan, and orange. Over the course of the experiment, the birds were presented with two keys, each illuminated with a different color. If a bird pecked the key illuminated with the higher-ranked color, then it was rewarded with its favorite food, pine nuts for pinyon jays and meal worms for scrub jays.

The experiment had two parts. In the first part, the birds were presented with only *adjacent* color pairs: red and green, green and blue, . . . , or cyan and orange. The birds were given daily sessions of 36 trials each, with the position of the higher-ranked stimulus randomized between left and right keys on each trial. And new color pairs were gradually added as the birds exhibited success in correctly selecting higher-ranked colors. Each of the birds was eventually required to track all six adjacent color pairs.

The pinyon jays reached a success rate of 0.80 in an average of 68 sessions, and after 100 sessions their accuracy was better than 0.85. The scrub jays learned significantly more slowly but eventually reached a similar level of accuracy.

In the second part of the experiment, the birds were also presented with *nonadjacent* color pairs. They were given 40 daily sessions of 36 trials each. During each session, they were presented with familiar, adjacent pairs of

3. The aim of the current article is to provide an explicit reinforcement model for the appropriation of a rule to a new context and discuss how such a model allows for the evolution of basic concepts. It also allows for a general evolutionary account of the *composition* of basic rules. The composition of rules is discussed briefly in Barrett (2013a) and is the focus of Barrett (2014a).

4. See also Barrett (2013a) for additional details regarding this and related experiments.

colors on 33 trials and novel, nonadjacent pairs on three trials. The empirical question for this part of the experiment was whether the birds would be able to determine the order of the nonadjacent color pairs on the basis of what they had learned from their experience with just the adjacent color pairs.

They were able to do so. Indeed, both species immediately exhibited a high level of accuracy on the trials involving the nonadjacent colors. The pinyon jays chose the correct color, as determined by the color order the experimenters initially assigned to the bird, on the nonadjacent pairs with an accuracy of 0.86, and for the scrub jays 0.77. The experimenters concluded that the birds were making transitive inferences on the basis of prior experience. But the birds were doing more than that.

A nonadjacent color judgment was taken to be correct in the experiment if and only if it agreed with the linear ordering the experimenters initially assigned to the bird, but the pairwise relation between adjacent colors that the birds learned in the first part of the experiment when presented with only adjacent colors does not by itself determine any relation whatsoever over the nonadjacent colors. Hence, the birds were both constructing a full linear color order from the partial information provided by just the nonadjacent color pairs and appropriating a previously acquired rule to make transitive inferences on the basis of this full linear order.

q2

3. A Simple Sender-Predictor Game. We consider the evolution of transitive rule-following and its subsequent appropriation to a new task in two steps: first, we consider how an agent might evolve a rule for transitive inference, then we consider how the agent might evolve to apply the old evolved rule to a new context. We start with a simple sender-predictor game.

In the simplest sort of sender-predictor game, a variety of the Skyrms-Lewis signaling game, there are two agents, a sender and a receiver who makes predictions on the basis of the sender's signal. The sender observes the state of nature, then sends a signal. The receiver, who cannot observe nature directly, performs a predictive action on the basis of the signal that either matches the future state of nature and is successful or does not match the future state of nature and is unsuccessful. If the act is successful, then the disposition that led to each agent's last action is reinforced; otherwise, it is not reinforced and may be weakened.

All it means for an action to be successful in the context of such a game is that it generates a result that, given the nature of the world that the agents inhabit and the agents' second-order dispositions to update their first-order dispositions to signal and to act, leads to the reinforcement of those first-order dispositions that produced the action; similarly, all it means for an action to be unsuccessful is that, given the nature of the world and the agents'

second-order dispositions, it generates a result that does not lead to the reinforcement of the first-order dispositions that produced the action.⁵

The second-order dispositions of the agents determine what resources they start with and how they learn using these resources. As a concrete example, consider four equally likely states of nature (0) sun, (1) fog, (2) rain, and (3) snow; four possible signals A, B, C, and D; and four predictive actions (0) bring sunglasses, (1) bring flashlights, (2) bring umbrellas, and (3) bring snowshoes. Further, suppose that the agents learn by bounded reinforcement with punishment.

More specifically, suppose that the sender has an urn corresponding to each of the four states of nature. Each urn contains balls corresponding to each of the four signal types. The receiver has an urn corresponding to each signal type, and each of these urns contains balls corresponding to each possible action (see fig. 1). The sender observes the state of nature, draws a random ball from the corresponding urn, then sends the signal indicated by the ball. The receiver sees the signal, draws a ball from the corresponding urn, then performs the action indicated by the ball. If the action is successful—if the action matches the future state of nature given the agents' second-order dispositions—then agents put the ball back into the urn they drew it from and add a ball of the same type unless there are already 1,000 balls of that type in the urn. If the action is unsuccessful—if the action does not match the future state of nature—then the agents do not return the ball they drew to the urn they drew it from unless it was the last ball of its type, in which case they simply return that ball.⁶

The agents' signals are meaningless when they begin to play this game, but as the sender's dispositions to signal (conditional on the state of nature) and the receiver's dispositions to act (conditional on the sender's signal) evolve, the sender's signals become meaningful in the precise sense of serving as a reliable basis for successful coordinated action. More specifically, on this evolutionary game, the sender and receiver start off randomly signaling and randomly acting. But, as they learn from experience, the sender and receiver typically (0.993) evolve a set of nearly optimal (0.994) dispositions (on 1,000 runs with 1×10^6 plays on each run). A slightly more

5. For discussions of Skyrms-Lewis signaling games and sender-predictor variants, see, e.g., Lewis (1969), Skyrms (2006, 2010), Barrett (2007, 2009, 2013b, and 2014b), and Argiento et al. (2009). See also Skyrms (2000) for an early model for the evolution of inferential behavior in a signaling game. While we are considering evolution in the context of a learning dynamics, there are population models that exhibit analogous features.

6. There are a number of different learning dynamics one might consider. See Roth and Erev (1995) for a discussion of other closely related options.




Figure 1.

subtle game allows for the coevolution of both a representation of states and transitive rule-following behavior.

4. A Model for the Evolution of a Transitive Rule. Consider a signaling game involving two senders, a left coder and a right coder, and an actor (see fig. 2). Together, the two coders and the actor might be thought of as the internal functional components of the representational system of a single agent. Two colors are randomly selected from a preordered set of seven colors with each pair of colors equally likely. One color is presented to the left coder, and the other is presented to the right coder. A play of the game is successful if the actor makes an order judgment, say thumbs-up ($a > b$), thumbs-down ($a < b$), or palm-down ($a = b$), that corresponds to the order in nature of the colors presented to the coders.

In this model, we consider a more sophisticated variety of learning than in the last.⁷ Each agent's urn begins with just a single black ball. Each coder draws a ball at random from the urn that corresponds to the color of its respective stimulus. If the ball is black, a new signal type is invented and sent to the actor; otherwise, a signal of the type of the drawn ball is sent to the actor. The actor has an urn corresponding to each pair of signals the coders might send. Each of these urns begins with a single ball of each action type: $a > b$, $a < b$, and $a = b$. If successful, the ball drawn from each urn is returned, and a new ball of that signal or act type is added to the urn; otherwise, the ball drawn from each urn is just returned. Finally, newly invented signal types are only kept if they lead to a successful action the first time they are used.⁸

7. See Argiento et al. (2009) and Alexander, Skyrms, and Zabell (2012) for discussions of this invention-learning rule and its properties.

8. This game, like the last, presupposes that agents can identify different states of nature and different signals. See Jäger (2007) and O'Connor (2014) for recent discussions regarding how categories might evolve. The current story might then be understood as explaining how such evolved categories might subsequently be employed to evolve basic concepts associated with a particular general rule.

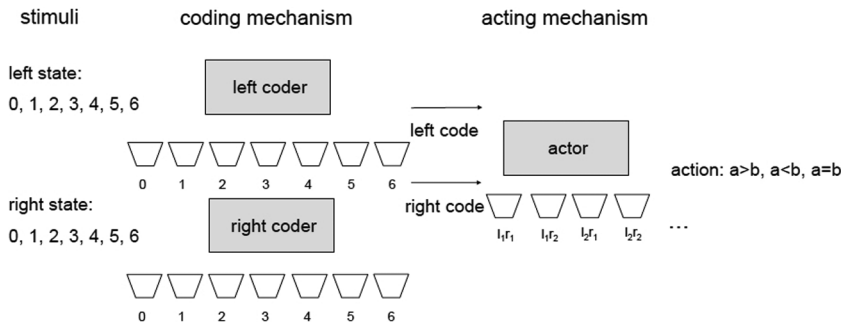


Figure 2.

Here the coders start by inventing new signals at a relatively high rate. They initially send these newly minted signals at random, and the actor initially acts randomly. But the composite system typically evolves to produce nearly uniform successful action. Specifically, after 1×10^7 plays, the cumulative success rate is typically (0.99) better than 0.75, and, in general, the more plays, the better the cumulative success rate. If one only requires that the system evolve the distinction between $a \geq b$ and $a < b$, then after 1×10^7 plays the cumulative success rate is typically (0.97) over 0.80, which is approximately the same accuracy exhibited by the jays.⁹ In each case, the system *invents* an internal language that represents the possible states of nature and *coevolves* a rule that linearly orders these states.

The representational system that evolves when the game is played with the full set of color pairs illustrates how it is possible for an agent to evolve an internal representation of the linear color ordering coded for in nature and the second-order dispositions of the agents with relatively modest evolutionary resources. In the model just described, the composite agent's evolved color-ordering dispositions constitute a rule that takes stimuli as input, represents the stimuli in an internal language, then outputs an action. Now we consider how, if these evolved color-ordering dispositions are triggered by a new type of experience, the composite agent may evolve to appropriate the old rule to structure the new experience.

5. A Model for the Evolution of Appropriation. Suppose that a coder-actor system has evolved a rule to represent and reliably judge the transitive

9. The cumulative success rate is a function of both how successful the process is and how quickly it gets there. It is a blunt but useful measure for comparing evolutionary processes—in this case, the process described here that evolves the transitive rule and the process described below that appropriates that rule to a new context.

relations of colors as described in the last section. Then consider presenting the old evolved system with a new set of stimuli that the composite system must learn to linearly order for successful action, say a set of seven musical tones.

The question then is whether the old system that evolved to order colors can be appropriated to the new task of ordering tones. The short answer is that it can if and only if the composite system can come to associate the new stimuli appropriately with the old inputs to the old evolved ordering system. And, it turns out, this can be accomplished by reinforcement learning on the new stimuli. Telling this part of the story provides a model for the jays' behavior when they appropriate an old rule of transitive inference to the new task of ordering colors and inferring their relations.

The old color-ordering system is represented by the three boxes on the right in figure 3. Each old coder urn represents a color, and we are supposing that the old system has evolved to correctly order color stimuli. The urns to the left correspond to the new tone stimuli. Each of the tone urns contains a ball for each of the old color urns. When each new coder gets a tone from nature, the coder draws a ball from the corresponding tone urn. This ball tells the coder which old color urn to then draw from. The coder draws a ball from the indicated color urn, then sends the indicated signal to the actor. Since the old color-ordering system has evolved to correctly order colors, the actor orders the signals as if the coders had observed the colors corresponding to the color urns they drew from. Consequently, optimally successful tone judgments will evolve if and only if the coders evolve to associate tones to the corresponding color urns.

If the coders had direct access to whether they had in fact chosen the color urns that corresponded (in the color ordering) to their tones (in the tone ordering), then they might very quickly and easily evolve a successful map from tones to colors.¹⁰ A more plausible evolutionary story, however, is one in which the only evidence the coders have concerning whether they have the right map from tones to colors is the order of judgments the actor in fact makes when the tones are treated as colors on each play of the game.

Suppose that the coders learn by simple reinforcement with punishment on the results of the actor's judgments. More specifically, on a play of the game, if the new coders choose balls that indicate color urns that in turn lead the actor to correctly order the tones, then each new coder returns her ball to the tone urn from which it was drawn and adds a copy of the same ball type; otherwise, each new coder discards the ball she drew unless it was the last

10. Indeed, in this case, on 1,000 runs, the composite system, using only simple reinforcement, evolves to successfully match the new tone stimuli to the corresponding old color-ordering system with an accuracy better than 0.80 on just 1×10^4 plays per run.




Figure 3.

ball of its type in the urn, in which case, she simply returns it to the urn. We suppose that the contents of the old color-ordering urns do not change on plays of this game.¹¹

The composite system typically (0.995) evolves to successfully match the new tone stimuli to the corresponding old color-ordering system with an accuracy better than 0.80 with 1,000 runs and 1×10^5 plays per run. The appropriation of the old evolved rule is much more efficient than evolving a new rule from scratch. In this case, evolving to appropriate the old ordering system to ordering tones is better than two orders of magnitude faster than the initial evolution of the color-ordering system.¹²

The evolutionary efficiency here comes from the actor's dispositions already being well tuned to making successful ordering judgments with the old stimuli. All that needs to be negotiated, then, is the map from the old to the new stimuli.

When the evolutionary process is successful, the system evolves to treat each tone as if it were the corresponding color in the color ordering induced by the old transitive rule. The coders then have evolved an analogy between tones and colors. And as other stimuli are similarly associated, the linear order represented in the dispositions of the composite system evolve to represent a linearly order over ever richer equivalence classes of stimuli. These equivalence classes might be taken to represent basic ordinal concepts under an increasingly abstract general rule that allows for transitive inference for whatever stimuli the system has in fact evolved to associate with these concepts.

11. The dispositions representing the old rule need not be fixed, but if they do evolve, they need to do so significantly more slowly than the process that appropriates the old rule to the new context.

12. It also involves a less sophisticated learning dynamics. Salient to this point, adding punishment to invention in the model that evolves the initial color ordering speeds that evolution, but it is still two orders of magnitude slower than the evolution of the map from tones to colors here. See Barrett (2013a) for a discussion.

6. Appropriation on Incomplete Evidence. One of the striking things about the behavior of the jays was that when presented with only adjacent color pairs they applied a previously acquired transitive rule to impose a linear order on nonadjacent color pairs as well. One sees a similar phenomenon in the current model.

When trained on a complete, unbiased set of tone stimuli, the old transitive rule is nearly always appropriated to the new context in such a way that the composite system is almost always successful on both adjacent and nonadjacent tone judgments. Specifically, on the learning dynamics described in the last section, the composite system evolves to represent the new tone stimuli in such a way that it exhibits a mean success rate in making order judgments on nonadjacent stimuli of 0.993 on 1,000 runs with 1×10^5 plays per run.

When trained on just adjacent tones, the composite system typically (0.766) evolves to correctly judge the order of new tone stimuli with an accuracy of better than 0.80 on 1,000 runs with 1×10^5 plays per run. Much of this success is due to the fact that the composite system evolves to do very well in ordering the adjacent stimuli that it is trained on. But, like the jays, the old evolved rule also induces a strong bias in how the new nonadjacent stimuli are ordered. Specifically, the composite system evolves to make the correct order judgments on nonadjacent tones when trained on only adjacent tones with a mean success rate of 0.703 on 1,000 runs with 1×10^5 plays per run.

The detailed behavior of the model is subtle. When trained on only adjacent tones, the composite system sometimes evolves to associate the new tone stimuli with the old color stimuli in a one-to-one manner that produces nearly perfect order judgments on the new nonadjacent stimuli, just as if the system had been trained on a complete unbiased set of tone stimuli.¹³ More often, when trained on only adjacent stimuli, the new stimuli are linearly ordered by the old rule in contiguous patches. Tones 0–2 might, for example, be linearly ordered and associated with the full range of colors 0–6, tones 4–6 might be similarly ordered and also associated with the full range of colors 0–6, and tone 3 might not be strongly ordered at all. While sub-optimal, the composite system typically does very well in ordering adjacent tones and much better than chance in ordering nonadjacent tones on such a representation.¹⁴

13. Such optimal behavior is observed in approximately 1 in 50 runs of the game as described. Even fixing the right pairwise order of the new stimuli, the coders need to be lucky to get the relative positions of tones and their corresponding colors right early in the evolutionary process.

14. See Bond et al. (2003) for details regarding the relative reliability of different types of order judgments for the two species. Getting the linear order right in patches is perhaps most similar to the behavior exhibited by the scrub jays.

In short, then, sometimes the old evolved rule induces a full linear order on the new stimuli on the basis of fundamentally incomplete evidence, evidence that itself does not entail a linear order at all. And even when such incomplete evidence leads to a suboptimal ordering of the new stimuli, the suboptimal ordering nearly always exhibits a strong linear bias on simulation, and it is the old rule that does the work of ordering the new stimuli in each of the linearly ordered patches.

7. Rule-Following, Analogy, and Basic Concepts. We have seen how it is possible for simple rule-following behavior to evolve in a general sort of signaling game. We have also seen how an old rule may evolve to be appropriated to a new context and how such appropriation may be much more efficient than evolving a new context-specific rule.

On the reinforcement model described here, the agents appropriate the old rule to a new context by evolving an analogy between the old and the new objects of experience in which the composite system comes to treat the new objects of experience as if they were the old objects of experience with respect to the old rule. As the rule is appropriated to other contexts, new analogies are formed, and the application of the rule becomes increasingly general.

In some cases the appropriation of a rule to a new context evolves a one-to-one map from the new objects of experience to the old objects. This typically happens on the model described here when the color-ordering rule is appropriated to ordering tones by reinforcement with punishment on the full set of new tone stimuli. When this happens, the analogy that evolves between the old and the new objects of experience induces equivalence classes of objects. As the evolved rule evolves to be similarly appropriated to other contexts, it becomes increasingly general. As the elements of each equivalence class play the same role under the evolved rule, the equivalence classes of objects induced by the evolved analogies might be thought to represent basic concepts, where what concepts they are is determined by the increasingly general rule that relates the elements of the evolved equivalent classes. Here, the equivalence classes formed by these associations might be thought to represent small finite ordinals that are linearly ordered under an increasingly general transitive inference rule that itself evolved in the context of promoting successful action given the agents' second-order dispositions and the nature of the world they inhabit.¹⁵

15. In contrast, in the context of the game described in Barrett (2013b), the concepts would be of small finite cardinals under the evolved rule of cardinal addition.

REFERENCES

- Alexander, J. M., B. Skyrms, and S. Zabell. 2012. "Inventing New Signals." *Dynamic Games and Applications* 2:129–45.
- Argiento, Raffaele, Robin Pemantle, Brian Skyrms, and Stas Volkov. 2009. "Learning to Signal: Analysis of a Micro-Level Reinforcement Model." *Stochastic Processes and Their Applications* 119 (2): 373–90.
- Barrett, Jeffrey A. 2007. "Dynamic Partitioning and the Conventionality of Kinds." *Philosophy of Science* 74:527–46.
- . 2009. "Faithful Description and the Incommensurability of Evolved Languages." *Philosophical Studies* 147 (1): 123–37.
- . 2013a. "The Evolution of Simple Rule-Following." *Biological Theory* 8 (2): 142–50. doi:10.1007/s13752-013-0104-4.
- . 2013b. "On the Coevolution of Basic Arithmetic Language and Knowledge." *Erkenntnis* 78 (5): 1025–36. doi:10.1007/s10670-012-9398-z.
- . 2014a. "The Evolution, Appropriation, and Composition of Rules." *Synthese*, forthcoming.
- . 2014b. "On the Coevolution of Theory and Language and the Nature of Successful Inquiry." *Erkenntnis* 79 (4): 821–34. doi:10.1007/s10670-013-9466-z.
- Bond, Alan B., Alan C. Kamil, and Russell P. Balda. 2003. "Social Complexity and Transitive Inference in Corvids." *Animal Behaviour* 65:479–87.
- Jäger, Gerhard. 2007. "The Evolution of Convex Categories." *Linguistics and Philosophy* 30:551–64.
- Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Cambridge, MA: Harvard University Press.
- Lewis, David. 1969. *Convention*. Cambridge, MA: Harvard University Press.
- O'Connor, Cailin. 2014. "Evolving Perceptual Categories." *Philosophy of Science*, in this issue.
- Roth, Al, and Ido Erev. 1995. "Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Intermediate Term." *Games and Economic Behavior* 8:164–212.
- Skyrms, Brian. 2000. "Evolution of Inference." In *Dynamics of Human and Primate Societies*, ed. Tim Kohler and George Gumerman, 77–88. New York: Oxford University Press.
- . 2006. "Signals." *Philosophy of Science* 75 (5): 489–500.
- . 2010. *Signals Evolution, Learning, and Information*. New York: Oxford University Press.

QUERIES TO THE AUTHOR

q1. AU: Your article has been edited for grammar, clarity, consistency, and conformity to journal style. Please read the article to make sure that your meaning has been retained. Note that we may be unable to make revisions that conflict with journal style or create grammatical problems. Thank you.

q2. AU: Italics omitted here and throughout when used for emphasis yet meaning seems clear without them or when not at first use of key terms.

q3. AU: (1) Each figure must be mentioned at least once in text, so parenthetical mentions of figures 1 and 2 were added accordingly. Also, (2) while not required, figures typically have brief captions describing their content. Do you wish to add captions to figures 1, 2, or 3?