The Evolution of Simple Rule-Following

Jeffrey A. Barrett

Biological Theory

ISSN 1555-5542

Biol Theory
DOI 10.1007/s13752-013-0104-4

Springer

Biol Theory
DOI 10.1007/s13752-013-0104-4

THEMATIC ISSUE ARTICLE: STRATEGIC INTERACTION

The Evolution of Simple Rule-Following

Jeffrey A. Barrett

Received: 29 February 2012 / Accepted: 25 September 2012

Abstract We are concerned here with explaining how successful rule-following behavior might evolve and how an old evolved rule might come to be successfully used in a new context. Such rule-following behavior is illustrated in the transitive judgments of pinyon and scrub-jays (Bond et al., Anim Behav 65:479–487, 2003). We begin by considering how successful transitive rule-following behavior might evolve in the context of Skyrms–Lewis sender–receiver games (Lewis, *Convention*. Harvard University Press, Cambridge, 1969; Skyrms, Philos Sci 75:489–500, 2006). We then consider two ways that an agent might come to use an old evolved rule in a new context. The first involves the agent evolving successful dispositions for one concrete type of experience, then associating a new type of experience with the old evolved dispositions. The second involves the agent evolving dispositions that represent a general inferential schema, then composing these dispositions with others in a way that allows the agent to make inferences concerning a new concrete type of experience.

Keywords Evolutionary game theory · Rule-following · Skyrms–Lewis sender–receiver games

The Evolution of Rule-Following Behavior

We are concerned here with explaining how successful rule-following behavior might evolve and how an old evolved rule might come to be successfully used in a new context. To explain this, one must be able to say how an agent might evolve to implement a successful rule, how the agent might come to use the old rule in a new context, and how this novel use might lead to successful action. For the last step, the evolved rule must in fact be appropriate to the new context and the agent must come to use the rule in the right way.[1]

We will consider how a type of Skyrms–Lewis sender–receiver game might lead to the invention and evolution of a representational system that produces successful transitive rule-following behavior.[2] We will then consider two ways that such successful rule-following behavior might be

J. A. Barrett (✉)
Logic and Philosophy of Science, University of California, Irvine, Irvine, CA, USA
e-mail: j.barrett@uci.edu

J. A. Barrett
Zukunftskolleg, University of Konstanz, Konstanz, Germany

[1] A rule is understood here as being implemented in the dispositions of an agent. In particular, rule-following behavior consists in the agent taking stimuli as input and producing an action as output. Rule-following in this sense, then, is easy. In contrast with the traditional Kripke-Wittgenstein rule-following worries, we are not concerned here with what the agents themselves may or may not know concerning the rules they are in fact following (Kripke 1982). Rather, we are concerned with how the agents might evolve successful rule-following behavior and how such successful rule-following might come to be extended to novel contexts. One might think of the present discussion, then, as providing a basic externalist evolutionary account of successful rule-following behavior.

[2] Simple sender–receiver games were first used by David Lewis (Lewis 1969) to account for the formation of convention. Brian Skyrms later described how they might be formulated in an evolutionary context without any assumptions of ideal rationality or common knowledge (Skyrms 2006). The type of invention-learning dynamics we will use here was proposed by Skyrms (Skyrms 2010). It is an extension of the sort of simple reinforcement learning proposed by Roth and Erev (Roth and Erev 1995). Some of its formal properties are considered by Skyrms, Alexander, and Zabell (Alexander et al. 2011). The present paper is closely related to Brian Skyrms' evolutionary models for how basic logical inference might evolve (Skyrms 2000) and my evolutionary models for how basic arithmetic language and practice might coevolve (Barrett 2013a). See also Simon Huttegger's discussion of the evolution of meaning (Huttegger 2007).

extended to contexts different from those in which they initially evolved. *Substitution* involves an agent evolving dispositions that represent a linear structure on one concrete type of experience, then substituting a new type of experience into the old evolved dispositions. *Composition* involves an agent evolving dispositions that represent the general properties of a linear structure, then using these dispositions to make inferences concerning a new concrete type of experience by composing representational systems. The sort of evolved rule-following behavior produced by composition can be more sophisticated than that produced by substitution, but it also requires more sophisticated evolutionary resources.

We will begin by considering a sort of transitive rule-following behavior observed in nature.

The Transitive Behavior of Jays

Pinyon jays *Gymnorhinus cyanocephalus* and scrub-jays *Aphelocoma californica* are two of several species of animal that exhibit surprisingly subtle sorts of transitive rule-following behavior. The transitive behavior of the jays is illustrated in the results of experiments performed by Bond et al. (2003).

In these experiments, a set of seven stimulus colors (red, green, blue, magenta, yellow, cyan, and orange) were randomly arranged in a linear order by the experimenters and a different such ordered set of colors was assigned to each of five adult pinyon jays and five scrub-jays. For convienience, we will represent the colors in a particular linear order by the numbers [0,6]. The birds were then presented with two keys illuminated with a color. If the birds pecked the key illuminated with the higher-ranked color, then they were rewarded with their favorite food, pine nuts for pinyon jays and meal worms for scrub-jays.

In the first experiment, the birds were presented with just adjacent color pairs to determine whether they might learn this pairwise ordering relation. Each of the stimulus keys were illuminated with a pair of adjacent colors, and the birds were rewarded for choosing the higher-ranked color as determined by the prior linear ordering assigned to each bird. A bird choosing color 3, for example, was rewarded if it was presented with color 2 but not if it was presented with color 4. The bird got the food if it pecked the higher-ranked key three successive times within 60 seconds.

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. The color pairs were gradually introduced, and more pairs were added as the birds exhibited success in correctly selecting higher-ranked colors. The birds were eventually required to track all six color pairs, which were presented in equal numbers and in a randomized order each day. The pinyon jays reached 80 % correct on all pairs in an average of 68 sessions, and after 100 sessions were better than 85 % accurate. While the scrub-jays learned significantly more slowly, they eventually reached similar levels of performance.

In the second experiment, the birds were tested to see whether they might now be able to determine the order of nonadjacent color pairs on the basis of what they had learned from their experience with adjacent color pairs. To this end, the birds were occasionally presented with novel, nonadjacent color pairs intermixed with the familiar stimulus combinations. While the birds had seen and learned to make judgments involving color 2 and color 5 before, for example, these colors had never been presented together. If the birds were able to reliably choose the higher-ranked color in a novel pair, this would provide evidence that they were somehow representing the pairwise relation between adjacent colors in the context of a linear order over all the colors, then making reliable inferences concerning the properties of this representation.

The birds were given 40 daily sessions of 36 trials each. During each session, they were presented with familiar, adjacent pairs of colors on 33 trials and novel, nonadjacent pairs on three trials. Both species immediately showed significantly higher accuracy on the trials involving non-adjacent colors than expected by chance. For pinyon jays the proportion of correct choices on novel stimuli was 0.86, and for scrub-jays it was 0.77.[3]

It is important to note that the birds are involved here in more than just transitive inference. The pairwise relation between adjacent colors that the birds learn in the first experiment does not by itself determine any relation whatsoever over the full set of colors. The experimenters might, for example, have chosen to reward pecking 1 when paired with 3 in the second experiment and this would still be a fair game. Part of what is interesting in the behavior of the birds, then, is the representational bias they exhibit when they act as if they expect the overall relation to be a

[3] The difference in accuracy is not statistically significant here, but there was a difference in the competence exhibited by the two species that depend on the position of stimuli in the overall implicit ordering. While the pinyon jay performance was equally good for all color pairs, one scrub-jay, for example, exhibited 92 % accuracy in responding to the color pair (1,2) but dropped to well below chance (37 %) for the color pair (4,5). Further, while the highest-ranked stimulus in the pair did not matter to the accuracy of the pinyon jays, they were significantly slower in responding to pairs that were lower in the sequence; and while the scrub-jays exhibited a clear first-item effect on accuracy, they only exhibited slight latency effects for stimuli that were lower in the sequence. The experimenters took this as evidence that there is a basic difference in the way each species represents the implicit color order [(Bond et al. 2003), p. 484].

linear ordering. To be successful here the birds need both a preexisting representational bias and the ability to carry out transitive inferences with respect to this bias. And the evidence suggests that both the pinyon jays and scrub-jays represent the relation between adjacent colors as part of a linear order over all colors, then make reliable inferences on the basis of this transitive structure.[4]

If this is right, the evolution of the transitive rule-following behavior exhibited by the jays involves two steps. The birds start the experiment already equipped with a system for representing a linear ordering, then they learn to apply that representational system to the new concrete case of the colors. We will consider how one might model the evolution of such a representational system and two ways it might be adapted to a new context. The first step of the evolutionary story involves a type of Skyrms–Lewis sender–receiver game that we will, given the purposes at hand, call a representational game.

Basic Representational Systems and Substitution

In the simplest variety of Skyrms–Lewis sender–receiver game there are two agents: a sender and a receiver. The sender observes the state of nature, then sends a signal. The receiver, who cannot observe nature directly, performs an action on the basis of the signal that either matches the state of nature and is successful or does not match the 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. Indeed, 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 signals are meaningless when the agents begin to play such a 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 the sender's signal) evolve, the sender's signals become meaningful insofar as they serve as the basis for successful coordinated action. Whether an ideally successful signaling system evolves in such a game depends on the number of states of nature, the statistical distribution of states of nature, the agents' signaling resources, their initial first-order dispositions to signal and to act, and the second-order learning dynamics that update these initial dispositions, but meaningful language evolves in a very broad assortment of such games.[6]

Here we will consider a somewhat more subtle sort of representational game. This game involves a coding mechanism, consisting of a left coder and a right coder, and an actor. Together the left coder, the right coder, and the actor are the functional components of a potential representational system. These functional components might be thought of as all residing and interacting with each other within a single agent. Whether a representational system evolves and, if so, what system evolves depends on the game played between the functional parts. If a representational system does evolve, the agent will have evolved an internal representation that ties his experience to successful action.

The first game is played as follows. Two colors are randomly selected from a set of seven linearly ordered set of stimulus colors [0–6] with each pair of colors equally likely.[7] One color is presented to the left coder and the other color is presented to the right coder. Each coder has seven urns, one corresponding to each of the possible color stimuli. Each urn begins with only a single black ball. Each coder draws a ball at random from the urn that corresponds to the color of the stimulus presented to that coder. If the drawn ball is black, a new signal type is invented then sent to the actor; otherwise, a signal of the type represented on the drawn ball is sent to the actor. The actor has an urn corresponding to each pair of signals the two coders might

[4] The positional effect discussed in footnote 3 also suggests that the pinyon jays are better at getting a linear representation over all of the colors. This is part of the evidence that the different species use different representational strategies. See D'Amato and Columbo (1988), Terrace and McGonigle (1994), and Delius and Siemann (1998) for examples of alternative models for transitive behavior as exhibited by humans and other species.

[5] While useful, the distinction between the agents' first-order dispositions (their dispositions to signal and to act) and their second-order dispositions (the learning dynamics that updates their first-order dispositions) is at best a rough one. It begins to unravel if their second-order dispositions are themselves allowed to evolve in

Footnote 5 continued

response to the evolution of their first-order dispositions. Indeed, we will allow for precisely this when we discuss how an agent's second-order dispositions may evolve to fit old first-order dispositions to a new type of experience.

[6] For discussions of standard Skyrms–Lewis sender–receiver games and other variants see Lewis (1969), Skyrms (2006), Argiento et al. (2009), Barrett (2007, 2009, 2013b), and Skyrms (2010). While we will restrict our attention here to evolution in the context of learning models, one should expect similar results in the context of population models as there is, for example, a close formal relationship between evolution under reinforcement learning and evolution under the replicator dynamics.

[7] Note that here we are allowing for any pair of colors to be selected and even for a color to be paired with itself.

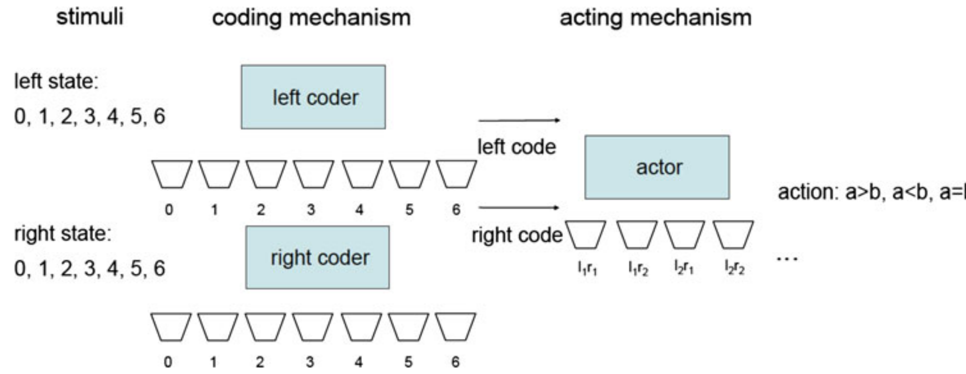

Fig. 1 A basic representational game with two stimuli

send. Each of these urns begins with a single ball of each action type $a > b$, $a < b$, and $a = b$. The actor draws a ball at random from the urn corresponding to the pair of signals sent by the coders and performs the corresponding action. To begin, we will suppose that this action involves the composite agent pecking one key $a > b$ or another $a < b$, or looking about in a confused way $a = b$. Here these are simply action types that are either successful or not depending on nature and the second-order dispositions of the functional parts of the potential representational system. We will later consider how such actions may themselves play the role of stimuli of types $a > b$, $a < b$, and $a = b$ to subsequent representational systems.

The dispositions of the functional components of the system are updated as follows. If the color presented to the left coder is higher-ranked than the color presented to the right coder and if the actor did action $a > b$ *or* if the color presented to the right coder is higher-ranked than the color presented to the left coder and if the actor did action $a < b$ *or* if the same colors are presented to both coders and if the actor did action $a = b$, then the ball drawn from each urn is returned and a new ball of the same signal or act type is added to that urn. Otherwise, the ball drawn from each urn is just returned to the urn from which it was drawn.[8] See Fig. 1 for an illustration of this game.

Such agents are easily simulated. When they begin, the coders do not have any signals that might represent the states of nature. They start by inventing new signals at a relatively high rate since there is not much else in their urns to keep them from drawing the black ball often. The coders initially send these newly-minted signals at random, and the actor initially acts randomly when he receives them. But as the first-order dispositions of the coders and actor are updated by reinforcement, the first-order dispositions of the parts of the system typically coevolve to represent the states of nature and the linear relationship between colors in the concrete sense that the actor does action $a > b$ when the left color is higher ranked than the right color, does action $a < b$ when the right color is higher ranked than the left color, and does action $a = b$ when the two colors are the same.[9] When such a representation of the colors and their linear ordering evolves, the coders' signals provide everything required for the actor to linearly order all possible pairs of the color stimuli. If only adjacent color pairs are presented to the coders, a representation evolves yet more quickly.[10] But in this case, of course, the color representation that evolves will only represent the order on adjacent pairs.

The representational system that evolves when the game is played with the full set of color pairs illustrates how it might be possible for an agent to evolve an internal representation of the linear color ordering in nature with relatively modest evolutionary resources. But this is not what the pinyon jays and scrub-jays are doing in the experiments described above. Rather, the birds begin with some sort of prior representation of a linear order, which they use to

[8] This dynamics describes a very simple sort of reinforcement learning with invention. Note that if a coder drew a black ball and if the action was successful, then the coder has invented a new type of signal he might send. The actor must then add urns corresponding to the new pairs of signals he may receive. Also, note that a new signal type is only kept if the first-play of the new type was successful, in which case, the new type becomes available for any representational purpose. As a result, the new signal type may well not evolve to mean what it meant on its first use.

[9] On simulation, after 1×10^7 plays, the cumulative success rate is better than 0.75 about 99 % of the time; and, in general terms, the more plays, the better the cumulative success rate on simulation. 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 over 0.80 about 97 % of the time, which is about the accuracy of the pinyon jays and scrub-jays. If one changes the learning dynamics to allow for *both* reinforcement on success and punishment on failure, then the convergence rates on simulation are orders of magnitude faster. In this sense, the simple reinforcement learning dynamics we are considering here might be thought of as difficult evolutionary context. The thought is that evolutionary stories that can be told here can be expected to be relatively robust.

[10] After 1×10^7 plays the cumulative success rate is over 0.90 about 85 % of the time.

impose structure on their new color experience and then to make inferences about the relation between colors that they have never before seen paired. That said, something akin to such rule-following behavior may occur if a new type of experience in fact produces signals in an old representational system. In particular, once a representational system for the full linear color ordering evolves, an agent may then exhibit a basic sort of rule-following behavior with respect to a new type of experience.

In the model just described, an agent's evolved color-ordering dispositions constitute a rule that takes stimuli as input and outputs an action.[11] If these color-ordering dispositions are triggered by a new type of experience, the agent will end up using the old rule to structure the new experience.[12] Here the agent's actions would order the new experiences as if they were colors.

The use of new stimuli to trigger old evolved dispositions may occur accidentally or as the result of an evolved or a learned association. But however it happens, the agent will end up structuring the new experience using a rule coded for in his old evolved dispositions. The moral is that all that is required for the agent to exhibit rule-following behavior in a new context is that his second-order dispositions allow a new type of input to trigger his old evolved first-order dispositions. If the old rule coded for in these dispositions is appropriate to the new context, the agent's actions on the new stimuli may be successful. But there can be no guarantee that the old rule will be appropriate to the new context. And even if the rule is appropriate to the new context, the agent may fail to coordinate the use of the rule to the context in a way that leads to successful action. While rule-following behavior just requires the evolution of consistent dispositions that tie experience to actions, successful rule-following behavior in a new context requires that one has evolved a rule that is in fact appropriate to the new context and that one in fact comes to use the the rule correctly.

Consider a concrete example. Suppose that an agent has evolved a linearly ordered system for the full set of colors above, then is presented with the tones G, B♭, C, D♭, D, F, and G′. If these tones trigger the color system, then the agent will act in such a way as to impose a linear order on the tones. But, of course, this does not mean that the agent's actions will be successful. Suppose that the agent is only successful in the new context if his actions order the tones by frequency. If the coders send the same signal for tone G that they send for color 0, the same signal for tone B♭ that they send for color 1, etc., then the representational system evolved for ordering colors would immediately allow for the reliable ordering by the frequency of the tones. But if the coordination between the old dispositions and the new type of experience were randomly determined, one should not expect successful action.

There are then two parts to a full story here. Once a representational system has evolved, one has evolved a rule that is coded for in the first-order dispositions of the agent and that might produce actions in a new context if the agent has second-order dispositions that match new states to old dispositions. But the agent will only be successful if the old rule is appropriate to the new context and is used in the right way. While the rule that orders colors might be used to order tones, one would like to be able to tell a plausible evolutionary story for how it might evolve to be used successfully. This would be an evolutionary story for how the second-order dispositions of the agent evolve to successfully coordinate the new sound stimuli with the old color-ordering dispositions. While the sender–receiver game described above does not have the resources to illustrate how second-order dispositions might evolve, once one has evolved a successful color representational system, it is not too difficult to tell a plausible evolutionary story for its successful appropriation.

Suppose that a successful full color-ordering system has evolved and that an agent begins by randomly pairing each tone with a color stimuli. From time to time a pair of tones is randomly selected and the colors that they are paired with are switched. If the switch yields a system that is more successful in ordering the tones, the new system is adopted; otherwise, the agent keeps the old system. On this simple learning dynamics, an agent typically finds an optimal pairing between tones and colors very quickly. Indeed, this evolutionary process is several orders of magnitude more efficient than the process that led to the initial color representational system.[13]

[11] For systems like those discussed here, this is typically a statistical rule that evolves to exhibit relatively stable behavior.

[12] It is assumed in the setup of a sender–receiver game that the senders have second-order dispositions that allow them to individuate types of states of nature and send signals on the basis of these state types. How their second-order dispositions individuate states of nature ultimately plays a role in determining what sort of signaling system they evolve. In this sense there is always an implicit similarity relation at work in the individuation of states of nature in a sender–receiver game. States of nature that are in fact grouped together under this relation will end up being treated the same way by the evolved representational system. Here we are allowing the second-order dispositions that represent this similarity relation to evolve.

[13] On simulation, the mean number of random switches that an agent must consider to get to the optimal pairing of tones and colors is about 40.82. Since choosing the most successful system is a matter of statistical inference, an agent following this learning dynamics faces a series of two-armed bandit problems. See Berry and Fristedt (1985) and, more recently, Zollman (2010) for discussions bandit problems and strategies for addressing them. In this particular case, the difference between the success rates of better and worse systems is relatively large, so while the agent may sometimes choose a worse system as better, it will not take much investigation to make the chance of doing so very small. More generally, however, the cost of comparing the success of alternative systems could make this learning dynamics less efficient.

That finding an optimal pairing between old dispositions and new stimuli can be significantly more efficient than evolving a new representational system may help to explain how an old system might come to be appropriated to a new context. The suggestion is that just as an agent might use a representational system that initially evolved to track colors to order tones, pinyon jays might use a representational system that initially evolved to track social order within a flock to order colors. When this happens, their newly learned second-order dispositions allow for the old representational system to serve as a general schema or rule that orders a new type of experience. In this way, one might understand the newly learned second-order dispositions to provide a form of analogical reasoning.[14]

Higher-Order Representational Systems and Composition

A more subtle sort of rule-following behavior might evolve in the context of a more sophisticated evolutionary model. We will briefly consider here how an agent might evolve dispositions that represent the general properties of a linear structure, then use these evolved dispositions to make inferences concerning a new type of experience. The idea is that a higher-order representational system might be composed with basic representational systems to produce actions that would not have been possible with the basic systems alone.

In addition to representing a more subtle sort of pattern, a higher-order system may take as inputs states resulting from its own actions or the actions of other representational systems.[15] This feature allows for simpler systems to act in sequence to form composite representational systems. And a composite representational system may exhibit new dispositions that result from the interactions between its parts. Here we are concerned with a higher-order system that may take its own actions or the actions of basic color systems as input.

An example of a higher-order transitive game is played as follows. There are two coders and one actor. Each coder is presented with a different aspect of the same state of nature. The left coder is sensitive to whether the state is of type $a > b$, $a < b$, or $a = b$ and the right coder is sensitive to whether the state is of type $b > c$, $b < c$, or $b = c$.[16] The states of nature are presented randomly, each equally likely. Each coder has three urns, one corresponding to each type of stimuli that they individuate, and each urn starts with a single black ball. The coders invent new signals and send them as before. The actor has an urn corresponding to each pair of signals the two coders may send. Here each of these urns begins with one ball of each action type $a > c$, $a < c$, $a = c$, and $a?c$. See Fig. 2 for an illustration of this game.

The dispositions of the functional components of the system are updated as follows. Let (s_l, s_r, d) be the situation where the state s_l is presented to the left coder, the state s_r is presented to the right coder, and the actor does action d. Suppose that nature and the second-order disposition of the agents are such that the signal types used by the agents are reinforced if and only if one of the following situations obtains:

$$\begin{array}{ccc} (a>b, b>c, a>c) & (a>b, b<c, a?c) & (a>b, b=c, a>c) \\ (a<b, b>c, a?c) & (a<b, b<c, a<c) & (a<b, b=c, a<c) \\ (a=b, b>c, a>c) & (a=b, b<c, a<c) & (a=b, b=c, a=c) \end{array} \quad (1)$$

Otherwise, the ball drawn from each urn is simply returned to the urn.

Such agents are easily simulated. Again, the coders start by inventing new signals quickly and using them at random and the resulting random actions are typically unsuccessful. Gradually, however, the coders typically evolve signals that represent states of types $a > b$, $a < b$, and $a = b$ and types $b > c$, $b < c$, and $b = c$ and the actor performs actions of types $a > c$, $a < c$, $a = c$, or $a?c$ as appropriate for inferences regarding the linear ordering represented in the payoffs of the game.[17]

[14] Such analogical behavior explains the differential behavior of the pinyon and scrub jays. In particular, Bond, Kamil, and Balda conclude that the social ordering that pinyon jays must navigate within a large stable flock explains their ability to get the full linear ordering right while scrub-jays, who typically live in pairs, are more likely to get later parts of the ordering wrong. More generally, they argue that animals living in large social groups should exhibit an enhanced capacity for transitive inference on partial evidence because they already require such inferential capacities for successful social interactions (Bond et al. 2003, pp. 479, 484–485).

[15] Since both evolve in the same general way, the distinction between basic and higher-order representational systems has more to do with what sort of pattern the systems end up representing and how the systems get used. While a basic system may evolve to represent a particular ordering among colors and take colors as inputs, a higher-order system may evolve to represent a property of an ordering generally and take actions of other representational systems as inputs.

[16] It does not matter what these input types are as long as (1) they in fact exhibit the general pattern of a linear order as described here and (2) the action of a basic color-ordering system produces a state that the higher-order system counts as one of these input types. This second condition allows the actions of a basic representational system to serve as a type of signal to a higher-order system. Note, however, that while the higher-order system must be able to take input from the actions of other systems, the significance of this input is not prearranged. Rather, in order for an agent to be successful, the coordination of the outputs of one system with the inputs of the next is something that one would expect to evolve as the result of the relative success of different ways of composing simpler representational systems to form more complex systems. We will discuss this further below.

[17] On simulation, after 1×10^7 plays, the cumulative success rate is better than 0.95 about 96 % of the time.

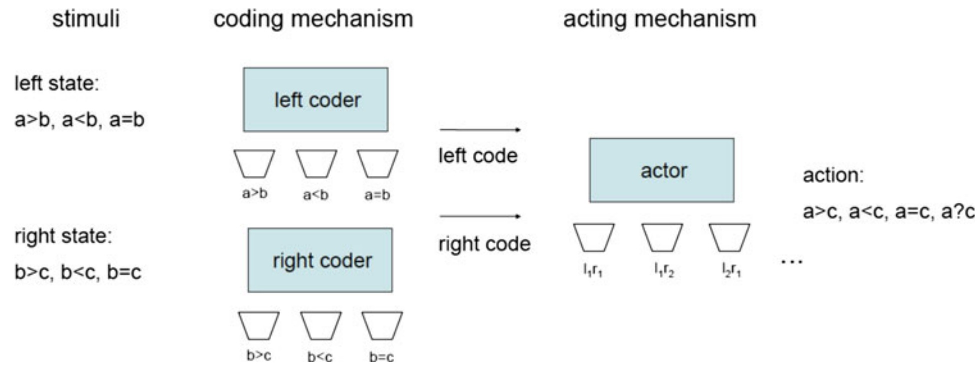

Fig. 2 A higher-order transitive game

We now have two types of representational system that might evolve by reinforcement learning with invention. We described earlier how a *basic* represenational system might evolve that takes two colors as inputs, codes these as signals, then outputs an action of type $a > b$, $a < b$, or $a = b$. For the purposes at hand, we will suppose that this basic system evolves in a context where the coders are presented with only adjacent and identical pairs of colors and where the system that evolves can, hence, only reliably judge the ordering of adjacent and identical pairs. This basic system represents, then, what the agent has learned about the ordering of adjacent colors. Here we have described how a *higher-order* representational system might evolve that takes a state of type $a > b$, $a < b$, or $a = b$ and type $b > c$, $b < c$, or $b = c$ as input, codes these types as signals, then outputs an action of type $a > c$, $a < c$, $a = c$, or $a?c$. This higher-order system evolves in a context where the coders are presented with states that in fact exhibit the structure of a linear ordering. The agent, then, coevolves a descriptive language for the salient features of such states and the ability to make reliable judgments concerning the transitive properties of this ordering on the basis of this representation. In this sense, the evolved higher-order system implements a general schema or rule for transitive inference. These two types of representational system are illustrated in Fig. 3.

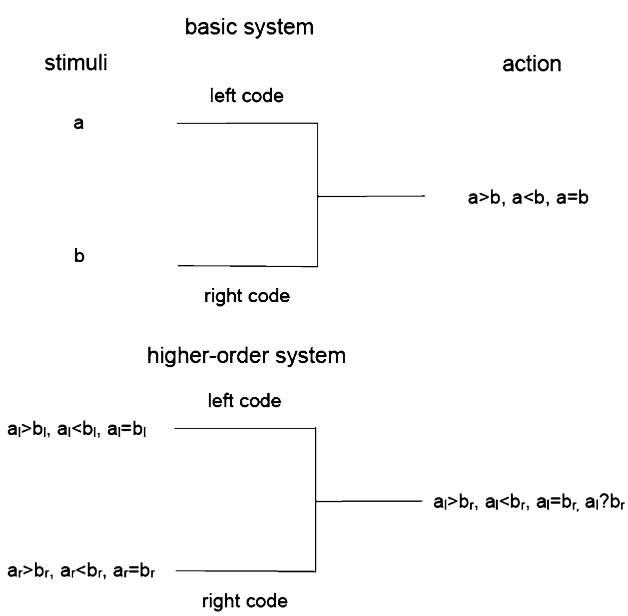

Fig. 3 A basic and a higher-order representational system

Consider the basic color-order system that evolves in a context where it is presented with only adjacent and identical colors and the higher-order transitive system just described. Again, the higher-order system may take as input its own actions and the actions of other representational systems, in particular, the basic color system. Suppose, then, that the actions produced by two applications of the basic color system are presented to the higher-order transitive system as illustrated in Fig. 4.

Suppose further that this composite system is applied to the adjacent color pairs (a_l, b_l) and (a_r, b_r) subject to the condition that $b_l = a_r$.[18] If the colors are in fact linearly ordered in nature, since each pair is made up of adjacent colors, the basic system can reliably judge their order, and the higher-order system can reliably judge the order of *nonadjacent* color pair (a_l, b_r). And recursive application of the two types of representational system would allow the agent to make reliable judgments regarding the relation between any pair of elements in the linear order from basic systems that represent only the relations between adjacent and identical pairs.[19]

More concretely, consider pinyon jays who have evolved a reliable higher-order transitive system then learn a basic color system for just the adjacent colors in the sequence [0,6] based on their experience in the laboratory. A pinyon jay's basic color system would produce an action of type $a_l > b_l$ when presented with stimulus (2,3) and produce an action of type $a_r > b_r$ when presented with stimulus (3,4) on a second application. Its higher-order

[18] Note that the condition $b_l = a_r$ can be reliably judged by agents who have evolved a basic system for adjacent colors as described above since they have also been trained up on identical colors.

[19] Note that the higher-order system would often be taking the results of its own actions as inputs on such a recursive application.

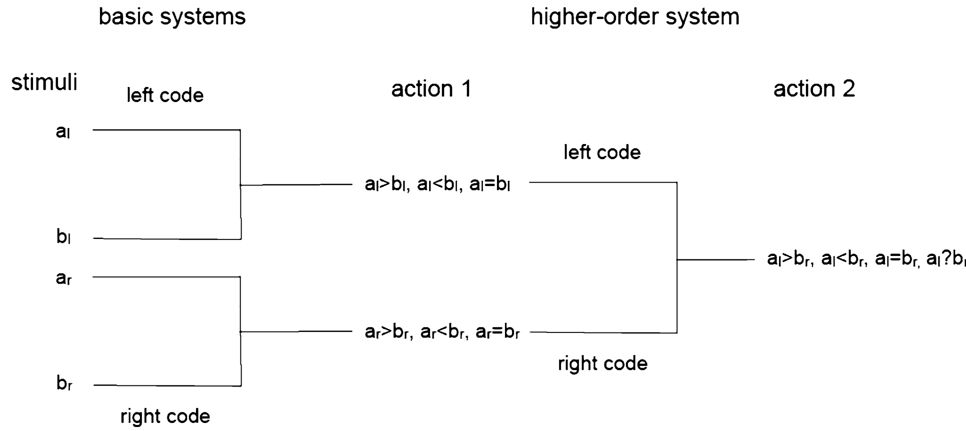

Fig. 4 A composite representational system

transitive system would then produce action $a_l > b_r$ when then presented with stimuli $a_l > b_l$ and $a_r > b_r$.[20] If the experimenters' prearranged relation between colors is in fact a linear ordering, this action reliably determines the order of the nonadjacent color pair (2,4). And recursive application of the procedure allows the jay to fill in the details of the prearranged order from the partial information provided by its experience of only adjacent pairs.

Of course, the use of this composite representational system might be disastrous if the relation rewarded by the experimenters were not a linear order. But, even if nature does reward the linear ordering associated with the order of the adjacent pairs, a composite representational system must be properly composed and constrained to impose the right relation on the complete set of colors. That is, the outputs from the basic color system must trigger the right inputs to the higher-order system and in the right way for the agent to impose the predetermined linear order on the color stimuli.

One would like to be able to tell an evolutionary story, then, for how a composite representational system constructed from pre-evolved subsystems might evolve to successfully represent a new pattern in nature. While the basic and higher-order sender–receiver games only explain how the subsystems might have evolved, one might begin to tell a story for how such subsystems might be used to construct composite systems that lead to successful action if one can settle on an appropriate space of the evolutionary options available. How difficult it would be for a successful composite representational system to evolve depends in part on what subsystems one has from prior evolutions and the degrees of freedom one has for composing them and regimenting their interactions. That said, there is good reason to believe that evolving successful new rule-following behavior from composition may be, in general, more difficult than evolving successful new rule-following behavior from substitution.

Consider the basic color system and the higher-order transitive system. Here the space of options one has for choosing three such systems, connecting them, and imposing identity constraints on their use is relatively small, and, hence, might be successfully searched by a learning or evolutionary dynamics that allows for mutation and preservation of the best structure tried so far.[21] But if one has more types of system to choose from or if one allows for larger composite systems or more complex coordinating conditions that regiment their interactions, one would expect the number of functionally distinct composite representational systems to grow very rapidly.

So, while having a higher-order system that represents a general schema for transitive inference might in principle be useful for agents who were repeatedly required to negotiate new linear orderings, it might in practice be difficult for an agent to fit such a rule to a new context without also having evolved auxiliary dispositions that help in composing representational systems and testing them for successful action.

[20] One might think of the actions produced by the basic systems as implemented by the jays as internal actions in thought. More specifically, one might imagine that the actions $a_l > b_l$ and $a_r > b_r$ of the basic color system produce brain records that are then read as input to the evolved higher-order inferential system. Of course, this recursive model may have precious little to do with actual pinyon jay cognition. Rather, it is just a how-possible explanation for the evolution of the sort of rule-following behaviors they exhibit.

[21] Setting aside the $a?c$ action, there are, on a relatively liberal notion of how one might compose subsystems and individuate functionally distinct composite systems, fourteen functionally distinct ways to arrange the basic and higher-order subsystems and six ways to associate the outputs of each of the two earlier subsystems with the inputs of a later system in the chain, for a total of 504 functionally distinct ways to compose the subsystems. In addition, there are fourteen functionally distinct ways to impose identity constraints on the initial inputs to the composite system. There are then 7,056 functionally distinct constructions and identity constraints to be explored by an evolutionary process.

Rule-Following and Rule-Following In a Novel Context

Sender–receiver games explain how representational systems might be invented and evolve and lead to successful action. Insofar as such systems implement rules that take experience and produce action, they also explain how the evolution of successful rule-following behavior is possible.

Successful rule-following behavior in the context that led to the evolution of the rule in the first place is automatic as long as nature is appropriately constant.[22] Substitution and composition explain how an old evolved rule may be used in a new context, but neither suggests that the resulting actions will be successful. For this, the old rule must be both appropriate to the new context and come to be used properly. While this might happen by chance, one would like to be able to tell a story for how this might be learned by the agent or happen by a plausible process of selection.

One can tell such stories for substitution and composition. In each case, the agent evolves a fit between his old representational systems and a new context by changing the second-order dispositions that determine how states connect to coder urns and how actor ball types connect to actions. As we saw in the case of fitting color judgments to tones, learning to fit an old evolved rule to a new context by substitution may be relatively easy. While composition is a more sophisticated learning strategy than substitution, and while there are obvious advantages to having a general representational system for transitive inference, for example, it may be significantly more difficult to fit an old evolved rule to a new context by composition depending on the number of degrees of freedom available for composing and regimenting subsystems.

The philosophical morals for rule-following are relatively straightforward. While one can readily explain how successful rule-following behavior might evolve in the context of a sender–receiver game, using an old evolved rule in new ways is always risky. An agent who does so always stands the chance of using the wrong rule for the task at hand or using the right rule in the wrong way. Only further experience can determine the appropriateness of the rule and fit it to the new context. And insofar as one has a satisfactory evolutionary story for how the old rule comes to be used in the new context, there is nothing mysterious in the agent's subsequent success in using it.

[22] Barrett (2013b) for an extended discussion of what aspects of nature matter for success in the sense of success appropriate to such evolutionary games.

Acknowledgments I would like to thank Brian Skyrms, Simon Huttegger, and Jason Alexander for helpful discussions and Martha Barrett for helpful comments on an earlier draft. This paper was written while visiting the Zukunftskolleg, Universität Konstanz. In this regard, I would also like to thank Franz Huber.

References

Alexander JM, Skyrms B, Zabell S (2011) Inventing new signals. Dyn Games Appl 2:129–145

Argiento R, Pemantle R, Skyrms B, Volkov S (2009) Learning to signal: analysis of a micro-level reinforcement model. Stoch Proc Appl 119:373–390

Barrett JA (2007) Dynamic partitioning and the conventionality of kinds. Philos Sci 74:527–546

Barrett JA (2009) Faithful description and the incommensurability of evolved languages. Philos Stud 147(1):123–137

Barrett JA (2013a) On the invention of numbers and the nature of mathematical knowledge. (forthcoming)

Barrett JA (2013b) On the coevolution of theory and language and the nature of successful inquiry. Erkenntnis. doi:10.1007/s10670-013-9466-z

Berry DA, Fristedt B (1985) Bandit problems: sequential allocation of experiments. Chapman and Hall, London

Bond AB, Kamil AC, Balda RP (2003) Social complexity and transitive inference in corvids. Anim Behav 65:479–487

D'Amato MR, Columbo M (1988) Representation of serial order in monkeys (*Cebus apella*). J Exp Psychol Anim Behav Process 14:131–139

Delius JD, Siemann M (1998) Transitive responding in animals and humans: exaptation rather than adaptation? Behav Process 42:107–137

Huttegger S (2007) Evolution and the explanation of meaning. Philos Sci 74:1–27

Kripke S (1982) Wittgenstein on rules and private language: an elementary exposition. Harvard University Press, Cambridge, MA

Lewis D (1969) Convention. Harvard University Press, Cambridge, MA

Roth Al, Erev I (1995) Learning in extensive form games: experimental data and simple dynamical models in the intermediate term. Games Econ Behav 8:164–212

Skyrms B (2000) Evolution of Inference. In: Tim K, George G (eds) Dynamics of human and primate societies. Oxford University Press, New York, pp 77–88

Skyrms B (2006) Signals. Philos Sci 75:489–500

Skyrms B (2010) Signals evolution, learning, & information. Oxford University Press, New York

Terrace HS, McGonigle B (1994) Memory and representation of serial order by children, monkeys, and pigeons. Curr Dir Psychol Sci 3:180–185

Zollman K (2010) The epistemic benefit of transient diversityy. Erkenntnis 72:17–35