THE EVOLUTION, APPROPRIATION, AND COMPOSITION OF RULES

JEFFREY A. BARRETT

Abstract. This paper concerns how rule-following behavior might evolve in the context of a variety of Skyrms-Lewis signaling game[13][15], how such rules might subsequently evolve to be used in new contexts, and how such appropriation allows for the composition of evolved rules. We will also consider how the composition of simpler rules to form more complex rules may be significantly more efficient than evolving the complex rules directly. And we will review an example of rule following by pinyon and scrub jays [11] as an illustration of the appropriation of a rule to a new context [4][3]. The proposal here is that the composition of rules might occur in a way that is precisely analogous to such simple appropriation. Finally, we will briefly consider how any finite truth-functional operation might evolve by the sequential appropriation of simpler rules.

1. Introduction

The present paper provides a model for the evolution of rule following behavior and for the evolution of more general, complex rules by the appropriation of more specific, basic rules. At the heart of the model are sender-predictor games, a variety of Skyrms-Lewis signaling game.

Sender-predictor games allow one to model the coevolution of simple representational systems and predictive rules [4][5]. Such rules may then be appropriated to contexts different from those in which they initially evolved and hence evolve from context specific rules to more abstract general rules [3]. We will be particularly concerned here with how appropriation allows for old rules to evolve to interact with implementations of themselves and other evolved rules. In particular, we will consider conditions under which the composition of old rules to form new rules may be significantly more efficient than evolving the new rules directly.

We will begin by discussing sender-predictor games, then consider how a two-sender game might evolve to implement a simple logical rule and review a model [4][3] for the transitive ordering behavior of pinyon and scrub jays [11] as an illustration of the modular appropriation of a new rule to an old context. The proposal is that the composition of evolved rules might occur in a precisely analogous way.

Finally, we will briefly consider how any finite truth-functional operation might evolve by the sequential appropriation of simpler rules.

2. Signaling games and the evolution of rule following

Signaling games were introduced by David Lewis [13] to provide a game-theoretic model for the establishment of convention. Signaling games were further developed as evolutionary games, both in learning and population contexts, by Brian Skyrms [15]. They show how simple signaling systems might evolve from initially random behavior. Here we will consider how the dispositions of agents evolve in the context of learning models.[1]

We will first consider a sender-predictor variant of the basic Skyrms-Lewis game. The suggestion will be that when such a game evolves successful action, the agents can be understood as having coevolved a simple representational system and a predictive rule.[2]

In the simplest sort of sender-predictor game, there are two agents, a sender who observers the state of nature and a receiver who makes predictions. 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 succeeds or fails to match the future state of nature. If the prediction is successful, then the disposition that led to each agent's last act is reinforced; otherwise, it is not reinforced and may be weakened.

What counts as success in such a model is a function of the dispositions of the agents and the nature of the world they inhabit. In short, a predictive action is successful if it leads to a result that, given the nature of the world that the agents inhabit and their second-order dispositions to update their first-order dispositions to signal and to act, leads to the reinforcement of those first-order dispositions to signal and act that ultimately produced the particular act. Similarly, a predictive action is unsuccessful if it generates a result that, given the nature of the world and the agents' second-order dispositions, does not lead to the reinforcement of the first-order dispositions that ultimately produced the act.

The second-order dispositions of the agents determine what learning resources they start with and how they learn using these resources. As a concrete example,

[1]While learning models and population models are closely related, it makes sense to choose one type of model and stick with it throughout the analysis. See [15] and [12] for discussions of the relationship between the evolution of signaling dispositions in the context of learning models and the evolution of signaling strategies in the context of the corresponding population models.

[2]See [4], [5] and [3] for discussions of this approach to the evolution of rule-following behavior. For further discussions of Skyrms-Lewis signaling games, sender-predictor variants, and alternative learning and evolutionary dynamics see [13], [17], [16], [9], [2], [8], [7], [6], and [15]. See also [10] in this issue for empirical evidence concerning how human agents in fact behave in the context of repeated signaling games.

suppose that the sender is sensitive to four equally likely and randomly distributed states of nature (0) *cool with a thick marine layer*, (1) *clear and crisp* (2) *dark clouds with an onshore flow*, and (3) *clear with Santa Ana winds*; that the sender has available four possible signals A, B, C, and D and that the receiver is sensitive to these signals; and that the receiver has available four possible actions (0) *bring jackets*, (1) *bring sunglasses*, (2) *bring umbrellas*, and (3) *bring sunglasses, sunscreen, and water*.

Concerning how the agents use these resources to update their first-order dispositions, in the context of the present game, and for all but one of the games that follow, we will suppose that the agents learn by bounded reinforcement with punishment.[3] More specifically, one might imagine that the sender has an urn corresponding to each of the four states of nature and that each urn contains balls corresponding to each of the four signal types. The receiver similarly has an urn corresponding to each signal type, and each of these urns contains balls corresponding to each possible act. The sender observes the early morning state of nature, draws a random ball from the corresponding urn, then sends the signal indicated by the ball. The receiver, who has no direct access to the morning state, sees the signal, draws a ball from the corresponding urn, then performs the predictive action indicated by the ball by bringing supplies for an afternoon trip to the beach. If the prediction is successful, if the receiver's action matches the afternoon state of nature, then each agent puts the ball back into the urn they drew it from and adds a ball of the same type unless there are already 1000 balls of that type in the urn. If the prediction is unsuccessful, if the receiver's action does not match the state of nature, then the agents do not return the balls they drew to the urns they drew them from unless a ball is the last of its type, in which case they simply return the ball.[4]

Concerning nature, we will suppose, for the sake of this particular game, that the weather is deterministic and such that a particular predictive action is always successful given the corresponding prior weather conditions.

The agents' signals are meaningless when they begin to play the game, but as the sender's dispositions to signal, conditional on the state of nature, and the receiver's dispositions to predict, conditional on the the sender's signal, evolve, the sender's signals come to serve as a reliable basis for successful coordinated action. While the sender and receiver start off randomly signaling and randomly acting, as they

[3]This is a learning dynamics that agrees well with human learning in a number of salient contexts. See [14] for a discussion of a closely related family of learning dynamics and relevant empirical evidence.

[4]Bounded reinforcement learning with punishment is much less sensitive to initial urn contents than, for example, simple reinforcement learning. That said, in the simulations here, each urn starts with one ball of each relevant type.

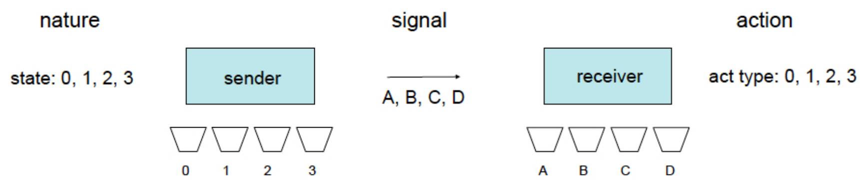

FIGURE 1. A simple signaling game

learn from experience, the composite system typically (0.993) evolves a set of nearly optimal (0.994) dispositions on 1000 runs with 10^6 plays on each run. When the two agents evolve to act together successfully, they have coevolved a representation of the relevant morning meteorological conditions and a rule for predicting the afternoon weather in Newport Beach on the basis of these conditions.

While it is often convenient to think of the sender and receiver as interacting agents, they might equally well be taken as representing different interacting functional parts of the same agent.[5] More specifically, the sender and predictor in the game above might be thought of as the perceptual system and deliberative system of a single agent respectively. On this interpretation, the game provides a very basic model for how an agent might coevolve an internal *representational system* for stimuli and a *predictive rule* that relies on that representational system.

The signaling games that follow are most naturally interpreted in as involving single agents.

3. The evolution of a simple truth-functional rule

Consider a game involving two senders, a left coder and a right coder, and a receiver, the actor. On each play of the game nature starts randomly in one of four possible states with equal likelihood. The states might be thought of as involving two possible conditions a and b, each of which either obtains or not. We will label the states, correspondingly, $0 : [0_a, 0_b]$, $1 : [0_a, 1_b]$, $2 : [1_a, 0_b]$, and $3 : [1_a, 1_b]$. To begin, we will suppose that the left and right coders each observe the full state of nature.

Each coder has four urns corresponding to each full state of nature labeled 0, 1, 2, and 3. Each urn contains balls labeled 0_l and 1_l for the left coder and 0_r and 1_r for the right coder. And each coder draws a ball at random from the appropriate urn and sends the corresponding signal to the actor. The actor does not know the state of nature, but it does know which coder sent each signal, and it has one urn for each possible pair of signals that it might receive from the two coders. There

[5]See [4] for a recent discussion of such an internalization of signaling games. Skyrms [16] has long considered such internal systems as interacting neurons as potentially implementing signaling games of one sort or another.

are two types of ball in each of the actor's urns: 0_c and 1_c. The actor draws a ball from the appropriate urn and performs the corresponding action.

As in the last game, we will suppose that the agents learn by bounded reinforcement with punishment. The actor's prediction is successful if $[0_a, 0_b, 1_c]$ (that is, if nature has the values 0_a and 0_b and the act is 1_c) or $[0_a, 1_b, 1_c]$ or $[1_a, 0_b, 1_c]$ or $[1_a, 1_b, 0_c]$; otherwise, the predictive act is unsuccessful. If the act is successful, each agent returns the ball to the urn form which it was drawn and adds a ball of the same type unless there are already 1000 balls of that type in the urn. If the act is unsuccessful, the agent does not return the ball to the urn from which it was drawn unless it was the last ball of its type, in which case the ball is simply returned to the urn.

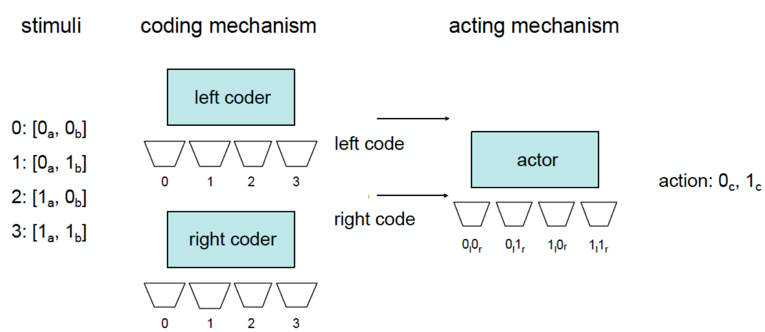

Figure 2. The evolution of **nand**

Again, the senders and receiver start off randomly signaling and randomly acting. On simulation, the composite system typically (0.907) evolves a cumulative success rate of better than 0.95 on 1000 runs with 10^6 plays each. Here, when the agents are successful, there is a sense in which the composite system implements the logical rule **nand** for truth values on the *ab*-partition. But the situation is subtle.

To be successful in this game, all the composite system has to do is to sort the four possible states of nature into those associated with output 0_c and those associated with output 1_c. And, even with just two signal types, a single sender has sufficient signaling resources to do this. The actor might then just attend to that sender's signals and act on how it sorts the states of nature. If so, while there would be a clear sense in which the composite system computes the operation **nand** for truth values under the *ab*-partition in this particular context, it would not implement **nand** as a truth-functional operation on two independent inputs.[6] A truth-functional rule that might be appropriated to other contexts is one where

[6]The effective **nand** rule that the composite system evolves and the truth-functional **nand** rule that one might have thought it evolved are different rules that happen to have the same extension in this particular context.

the receiver's act depends on the input to each sender and the signals each coder sends on this input. Such a rule is more likely to evolve when the senders each have access to different, independent aspects of nature.[7]

Consider the same game but where the left sender only has access to the truth value of a and the right sender only has access to the truth value of b. On simulation, the senders and actor sometimes 0.090 evolve a cumulative success rate of better than 0.99 and typically 0.720 evolve a evolve a cumulative success rate of better than 0.95 on 1000 runs with 10^4 plays each.[8] And here, when it does successfully evolve, the **nand** rule represents a truth-functional operation that depends on both inputs and can hence be appropriated to other contexts as a logical rule.

The appropriation of an old rule to a new context occurs when the old rule evolves to accept the new stimuli to accomplish a new, but typically similar, task. More expressively, the old **nand** rule is appropriated to a new context by evolving an analogy between the new stimuli and the stimuli responsible for the original evolution of the old rule. Once articulated to the old rule, the new stimuli lead the old actor to those actions it would have performed if the old rule were presented with the corresponding old stimuli.

In order to see how a rule might be appropriated to a context different from that in which it evolved, we will consider the evolution of a linear ordering rule, then the simple appropriation of that rule to a context involving new stimuli. In particular, we will consider an experiment described in [11] then review an evolutionary model developed in [4] to explain how such rule-following behavior might evolve then be appropriated to order new stimuli. We will then return to a discussion of **nand** to consider how such appropriation also allows for the modification and composition of evolved rules.

4. The evolution and appropriation of a transitive ordering rule

Pinyon jays *gymnorhinus cyanocephalus* and scrub jays *aphelocoma californica* are two of a number of species that have evolved to exhibit transitive rule-following behaviors. The rule following behavior of pinyon and scrub jays is illustrated in an experiment reported by Alan B. Bond, Alan C. Kamil, and Russell P. Balda.[11][9]

In the experiment, seven stimulus colors were arranged in a random linear order that was fixed for each bird. The birds were then presented with two keys, each

[7]In his model for the evolution of **xor**, Skyrms [17] gave each of the two senders access to only one aspect of nature. While this may at first appear to be an artificial constraint, it provides a context where the receiver must use information from both senders and hence allows for the evolution of a generalizable truth-functional operation.

[8]The composite system does better on longer runs. The results for runs of length 10^4 are mentioned here for the comparison of relative efficiencies later.

[9]This section is a short review of a longer discussion of this model in [3]. See [4] for another compositional model of the evolution of rule-following behavior.

illuminated with a different color. If a bird pecked the key illuminated with the higher-ranked color, then it was rewarded. The birds were initially presented with only *adjacent* color pairs: red and green, green and blue, . . . , or cyan and orange, and the position of the higher-ranked stimulus randomized between left and right keys on each trial. 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 learned to choose the higher-ranked color better than 0.85 of the time. The scrub-jays learned more slowly, but eventually reached a similar level of accuracy.

The birds were then presented with *nonadjacent* color pairs. The empirical question was whether the birds would be able to determine the order of the nonadjacent color pairs based on what they had learned from their experience with just the adjacent color pairs. Both species immediately exhibited a high level of accuracy on the trials involving the nonadjacent colors. The pinyon jays chose the higher-ranked color, as determined by the color order that the experimenters initially assigned to the bird, on the nonadjacent pairs with an accuracy of 0.86, and the scrub jays with an accuracy of 0.77. The experimenters concluded that the birds were making transitive inferences based on prior experience. But the birds were doing more.

A nonadjacent color judgment was taken to be correct in the experiment 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 when presented with only *adjacent* colors does not determine any relation whatsoever over the *nonadjacent* colors. Rather, to be successful in the task as presented, the birds were constructing a full linear color order from the partial information provided by just the nonadjacent color pairs then appropriating a previously acquired rule to make transitive inferences on the basis of this full linear order.

In order to model this behavior, consider another signaling game with a left coder, a right coder, and an actor. Here the two coders and the actor might be thought of as the functional parts of a single bird. In this game, 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, then they signal the actor. The actor performs one of three types of act: (0) $a > b$, (1) $a < b$, or (2) $a = b$. We will suppose that each such act does something distinctive that might subsequently be detected as the input to another process. This will be important for the discussion later. For present purposes, we will simply count an act as successful if it correctly represents the pre-specified order of the colors presented to the coders.

Here we will suppose that the composite system learns by reinforcement with invention.[10] Suppose that each coder'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.[11] Each of these urns begins with a single ball of each act 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. And newly invented signal types are only kept if they lead to successful act the first time they are used.

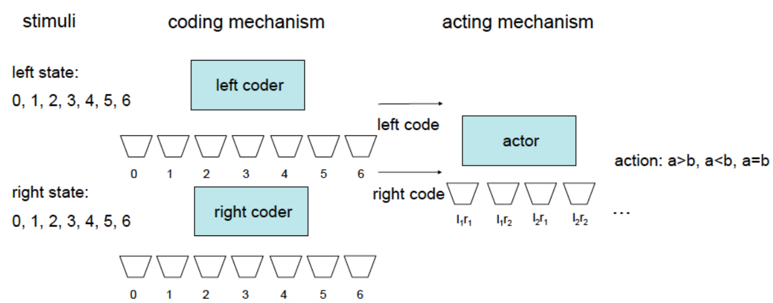

FIGURE 3. The evolution of an ordering rule

Here coders start by quickly inventing an assortment of new signals. They initially send these signals at random, and the actor acts randomly. After 10^7 plays, however, the cumulative success rate is typically (0.99) better than 0.75; and, in general, the more plays, the better the cumulative success rate.[12] When it evolves to be successful, the composite system has invented an internal language that represents the possible states of nature and has coevolved a rule that linearly orders these states.

The dispositions of the composite system constitute a rule that takes color stimuli as input, represents the colors as signals, then outputs an act that indicates the order of the colors. Once evolved, this rule might be appropriated to order a new type of stimuli in the context of a new task.

As an example such an appropriation, suppose that the composite system must learn to linearly order seven musical tones to allow for successful action. The old

[10] See [1] and [2] for discussions of this invention-learning rule and its properties.
[11] New urns are introduced for pairs of signals as new signals are invented by the coders.
[12] If one only requires that the system evolve the distinction between $a \geq b$ and $a < b$, then after 10^7 plays the cumulative success rate is typically (0.97) over 0.80, which is approximately the same accuracy exhibited by the jays.

color-ordering rule might be appropriated to the new task of ordering tones by evolving an analogy that associates the new tones in a on-to-one manner with the corresponding colors in the old linear order. If the process of evolving such an analogy is more efficient than the process of evolving a new rule from scratch, then the appropriation of old rule for the new task at hand might be evolutionarily favored. And if the acts of evolved rules do something that might subsequently be detected as the input to another rule, then the appropriation of an old rule to a new context will also allow for the composition of rules. We will return to these two points after considering how the simple appropriation of a rule might work.

Suppose that a color-ordering rule has successfully evolved and is represented by the composite system on the right in the figure below (the old coding and old acting mechanisms). Each old coder urn represents a color. The urns to the left (the new coding mechanism) correspond to the new tone stimuli. Each of the tone urns contains a ball for each of the old color urns. When a new coder gets a tone from nature, it draws a ball from the corresponding tone urn. The ball tells the coder which old color urn to draw from. The coder draws a ball from the indicated color urn, then sends the corresponding 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, successful tone judgments evolve if the new coders evolve to associate tones to the corresponding color urns. Importantly, we will suppose that the only evidence the new coders have concerning whether they have the right map from tones to colors is the order judgments that the actor in fact makes when tones are treated as colors on each play of the game.

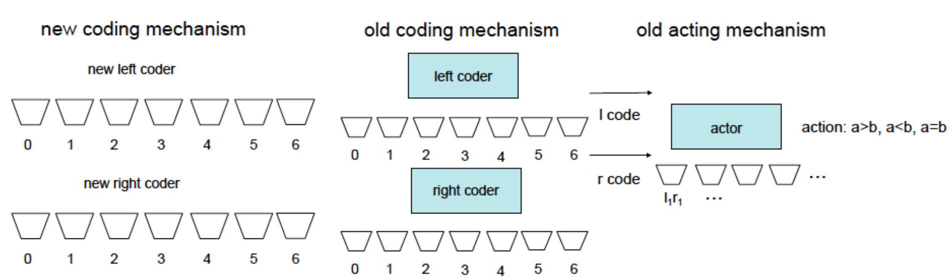

FIGURE 4. The simple appropriation of an evolved rule

We will suppose that the coders learn by simple reinforcement with punishment on the results of the actor's judgments. 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 ball of its type in the urn, in which case, she simply returns it to the urn. We will also suppose that the contents of the old color-ordering urns do not change on plays of this game. The dispositions representing the old rule need not be fixed, but if they do evolve, they must do so more slowly than the the process that appropriates the old rule to the new context.

Here 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 1000 runs and 10^5 plays per run. Note that the appropriation of the old evolved rule is significantly more efficient than evolving a new rule from scratch. Indeed, the appropriation of the old rule to the new context is better than two orders of magnitude faster than the initial evolution of the color-ordering system and involves a less sophisticated learning dynamics. The evolutionary efficiency comes from the actor's dispositions already being well-tuned to making successful ordering judgments with the old stimuli. The new coders just have to find an analogy between the old and new stimuli.

Salient to modeling the behavior of the jays, when the composite system evolves in the context of color ordering and the new coders are trained on just *adjacent* tones, the composite system evolves to judge the full order of tones with an accuracy better than 0.80. Here the new coders sometimes evolve to match tones to colors in a one-to-one manner that respects the full color ordering, but more typically they use the color ordering to group the tones into two or three linearly ordered segments. Within each segment, the composite system does very well in ordering both adjacent and nonadjacent tones. Between segments, the judgments are less reliable, but in aggregate, the composite system does about as well as the judgments of the birds trained on just adjacent colors.[13]

5. Modification of a rule by appropriation

When an old rule is appropriated to a new context, the new coders articulate new stimuli with the old rule to accomplish a new task. The new task might be nearly identical to the task that the old rule evolved to address as in the transitive ordering case just discussed. But it is also possible that, in addition to fitting the old rule to a new context, the appropriation of a rule significantly modifies what the old rule does. To see how the articulation of an old rule to a new context can significantly change the behavior of the old rule, we will briefly consider how an implementation of **nand** might evolve to implement **not**.

Consider an implementation of **nand** in a context with new left and right coders that each see the same two states of nature 0_a and 1_a. Each new coder has two

[13]See [3] for a detailed discussion of the appropriation of the old ordering rule on the incomplete evidence of only adjacent tones.

urns corresponding to the two states of nature and each urn contains balls that correspond to each of the four urns of the old left and right coders respectively. Each new coder observes the state of nature, then draws a ball from the corresponding urn. The drawn balls determine which of the old left- and right-coder urns is drawn from to determine the signals to the actor. The signals are sent, and the actor acts on them in the usual **nand** way. If the act is 1_c and the state of nature was 0_a or if the act is 0_c and the state of nature was 1_a, then the new coders return their balls to the urns from which they were drawn and, if there are fewer than 1000 balls of that type, add a ball of the same type; otherwise, they discard the ball they drew unless it was the last ball of its type in the urn, in which case, it is simply returned to the urn. We will suppose again that the dispositions represented in the old rule do not change on plays of the new game.

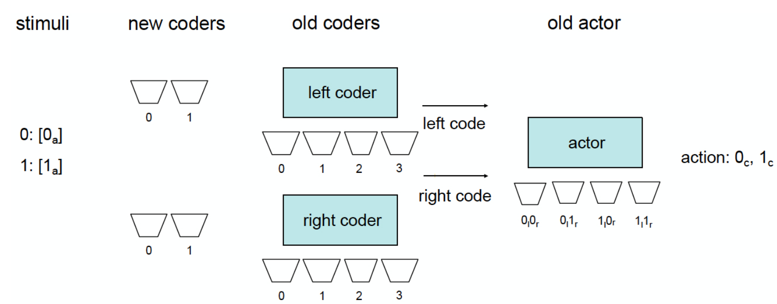

FIGURE 5. The appropriation of **nand** to evolve **not**

On simulation, the new coders typically 0.728 learn to articulate the new stimuli to the old **nand** rule to produce the **not** operation with a cumulative success rate of better than 0.99 on 1000 runs with 10^4 plays per run. In doing so, the new coders modify the effective behavior of the old **nand** rule dramatically.

There are, however, more efficient ways for **not** to evolve. Consider a one-sender game with two states of nature, two possible signals, and two possible acts and take a play of the game to be successful if and only if nature is 0_a and the actor does 1_c or if nature is 1_a and the actor does 0_c. For the purpose of comparison, suppose that that the sender and actor learn by bounded reinforcement with punishment precisely as described above. On simulation, this composite system typically 0.988 achieves a cumulative success rate of better than 0.99 on 1000 runs with 10^4 plays per run. So the appropriation of **nand** above is somewhat less efficient than the direct evolution of **not** here.

Whether evolving a new rule by the appropriation of an old rule is more efficient than evolving the new rule directly depends on details of the evolutionary context and the particular rules involved. One can often get a sense of the relative difficulty of the appropriation of an old rule to an new context and the direct evolution of the new rule by considering the degrees of freedom and the sizes of these degrees of freedom, as measured by the number and size of the urns involved, that must be coordinated for the successful evolution of the rule. Unsurprisingly, the overall similarity of the dispositions captured by the old and new rules also matters. In the present example, while **nand** is in some ways similar to **not**, one cannot take advantage of this similarity for evolutionary efficiency here since the appropriation of **nand** and the direct evolution of **not** each involve coordinating two degrees of freedom of roughly similar size.

That said, as we saw in the last section, evolving a rule for a new context by the appropriation of an old rule can be much more efficient than evolving the rule in the new context from scratch. This is particularly true if the old rule is being used to do something very like what it did in the context in which it initially evolved. We will consider another example of efficient appropriation in the next section.

6. The composition of rules by appropriation

Composing rules by appropriation to form a new rule is often more efficient than evolving the new rule directly. When an old rule is appropriated to a new context, the stimuli to which the new coders are sensitive might be aspects of nature directly or outputs from other evolved rules. The suggestion here is that the composition of rules works in precisely the same way as the simple appropriation of a rule except that at least part of the input to the old rule is articulated with the output of other evolved rules.

Consider the evolution of logical **or** from the composition of **not** and **nand**. Just as the ordering rule above may be appropriated to treat tones like colors, an implementation of **nand** may be appropriated to treat the output from two implementations of **not** like inputs from nature directly and, in the process, evolve to behave like **or**. Given two implementations of **not** that accept inputs from nature, the evolution of **or** by the appropriation of **nand** can be significantly more efficient than the direct evolution of **or**.

Suppose that one implementation of **not** has evolved to compute the negation of the truth value of a and another implementation of **not** has evolved to compute the negation of the truth value of b. Suppose further that one new coder is sensitive to $\neg a$ from the first implementation of **not** and that another new coder is sensitive to $\neg b$ from the second implementation. The new coders (represented by the stars in the figure below) have the evolutionary task of articulating these outputs to the old **nand** rule (represented by the dispositions of the system in the dotted box) to produce $a \vee b$ as its acts. If successful, the composite system implements **or** on the inputs a and b from nature.

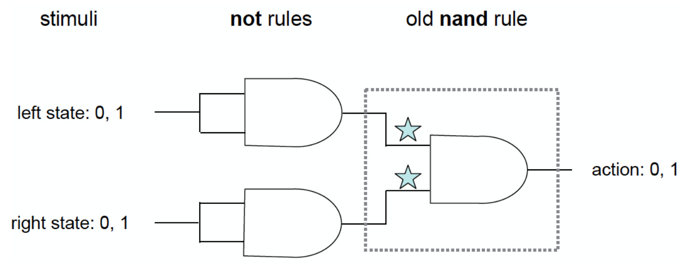

FIGURE 6. The evolution of logical **or** by composition

On simulation, the composite system typically 0.651 achieves a cumulative success rate of better than 0.99 on 1000 runs with 10^4 plays per run, and 0.959 of the time it does better than a cumulative success rate of 0.950. Similar to the direct evolution of **nand** described earlier, when **or** is evolved directly with bounded reinforcement with punishment, the composite system only sometimes 0.093 evolves a cumulative success rate of better than 0.99 and typically 0.756 evolves a cumulative success rate of better than 0.950 on 1000 runs with 10^4 plays each. The upshot is that the evolution of **or** by the appropriation of **nand** is significantly more efficient than the direct evolution of **or**.

The evolutionary efficiency here comes from the fact that **nand** is well-suited to representing logical **or** in this context and the appropriation of **nand** involves fewer degrees of freedom of similar size, measured in number and size of the urns involved, than the direct evolution of **or**. The new coders just need to learn how to map the outputs from the two implementations of **not** to the inputs of **nand**, which is a relatively easy task since it involves fewer degrees of freedom than involving **or** directly.

More generally, the stepwise evolution of a complex rule might be modeled as the sequential articulation of old rules, one at a time, with already situated rules. At each step, the new coders just have to fit the inputs to the next basic rule with the outputs from the situated rules. When there are appropriate basic rules available to compose, such a process can be significantly more efficient than evolving the complex rule directly. If the available basic rules are not perfectly well suited to the task at hand, as we saw in the last section, the appropriation of the basic rule to the new context may also tune that rule to the task at hand. But then the evolutionary task faced by the new coders can be significantly more difficult.

Since any finite truth-functional operator can be represented by a finite combination of implementations of **nand**, any finite truth-functional operator might evolve by a stepwise sequence of appropriations. One might think of the resulting complex rule as consisting in implementations of **nand** glued together by new coders that have evolved to articulate these implementations to the task at hand. Insofar as such stepwise appropriations may be more efficient than evolving the complex rule directly, evolution by modular appropriation may be evolutionarily favored in some contexts.

In the examples we have considered here, we have held the old evolved rules fixed on appropriation. One should, however, expect the old rules to continue to evolve after being appropriated. The model described here requires that the old rules evolve more slowly then the new coders, but it also allows for the continued evolution of the old rules. Since the evolution of the old rule in the new context involves more degrees of freedom than appropriation, this may provide a mechanism for better fit and perhaps better overall accuracy of the composite rule than would be possible by just the evolution of the new coders alone.[14]

7. Conclusion

We have considered how simple rule-following behavior might evolve in the context of signaling games, how such rules may then evolve to be appropriated to contexts different from those where they initially evolved, and how the modular appropriation of a rule allows for the composition of rules. In particular, the present model explains how it is possible for more complex rules to evolve from simpler parts, and, concretely, how any finite truth-functional operation might evolve stepwise by the sequential appropriation of simpler rules. In brief, the composition of rules works in precisely the same way as the simple appropriation of the ordering rule in the evolutionary model of the pinyon and scrub jays. Since the stepwise appropriation of simpler rules to evolve more complex rules takes advantage of the evolved dispositions of the actors, it often involves fewer degrees of freedom than evolving the complex rule directly. When it does, evolution by the composition of basic rules may be more efficient than the direct evolution of the complex rule.

References

[1] Alexander, J. M., Brian Skyrms, Sandy Zabell (2011) "Inventing New Signals" *Dynamic Games and Applications* (2011): 1-17.

[2] Argiento, Raffaele, Robin Pemantle, Brian Skyrms and Stas Volkov "Learning to Signal: Analysis of a Micro-Level Reinforcement Model," *Stochastic Processes and their Applications* 119(2):373–390.

[3] Barrett, Jeffrey A. (2014) "Rule Following and the Evolution of Basic Concepts" forthcoming in *Philosophy of Science*.

[14] The present model of composition by appropriation presents a number of questions. Since one would expect error to be cumulative as old rules are composed, one might wonder whether is there a plausible mechanism by which the old rules might be more precisely tuned to the task at hand once they are situated by the articular of the new coders. One mechanism that might serve this purpose is the continued evolution of the new coders once an old rule is appropriated to the new context. Another is the continued evolution of the old rules themselves.

[4] Barrett, Jeffrey A. (2013a) "The Evolution of Simple Rule-Following" *Biological Theory*. doi:10.1007/s13752-013-0104-4.

[5] Barrett, Jeffrey A. (2013b) "On The Coevolution of Basic Arithmetic Language and Knowledge" *Erkenninis*. doi:10.1007/s10670-012-9398-z.

[6] Barrett, Jeffrey A. (2013c) "On the Coevolution of Theory and Language and the Nature of Successful Inquiry" *Erkenntnis*. doi:10.1007/s10670-013-9466-z.

[7] Barrett, Jeffrey A. (2009) "Faithful Description and the Incommensurability of Evolved Languages," *Philosophical Studies* 147(1): 123–137.

[8] Barrett, Jeffrey A. (2007) "Dynamic Partitioning and the Conventionality of Kinds," *Philosophy of Science* 74:527–546.

[9] Barrett, Jeffrey A. (2006) "Numerical Simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar," *Institute for Mathematical Behavioral Sciences*: Paper 54. http://repositories.cdlib.org/imbs/54

[10] Bruner, Justin, Simon Huttegger, Cailin O'Connor, and Hannah Rubin. "David Lewis in the Lab: An Experimental Study of Signaling Conventions" manuscript.

[11] Bond, Alan B., Alan C. Kamil, and Russell P. Balda (2003) "Social Complexity and Transitive Inference in Corvids," *Animal Behaviour* 65: 479–487.

[12] Huttegger, Simon and Kevin J. S. Zollman (2011) "Signaling Games: Dynamics of Evolution and Learning," in *Language, Games, and Evolution* Trends in Current Research on Language and Game Theory Series: Lecture Notes in Computer Science, Vol. 6207, A. Benz, C. Ebert, G. Jäger, R. van Rooij (eds.), Berlin Heidelberg: Springer.

[13] Lewis, David (1969) *Convention.* Cambridge, MA: Harvard University Press.

[14] 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.

[15] Skyrms, Brian (2010) *Signals Evolution, Learning, & Information*, New York: Oxford University Press.

[16] Skyrms, Brian (2006) "Signals," *Philosophy of Science* 75(5): 489–500.

[17] Skyrms, Brian (2000) "Evolution of Inference," in Tim Kohler and George Gumerman (eds.) *Dynamics of Human and Primate Societies*, New York: Oxford University Press, 77–88.