

Self-assembling Games

Jeffrey A. Barrett and Brian Skyrms

ABSTRACT

We consider how cue-reading, sensory-manipulation, and signalling games may initially evolve from ritualized decisions and how more complex games may evolve from simpler games by polymerization, template transfer, and modular composition. Modular composition is a process that combines simpler games into more complex games. Template transfer, a process by which a game is appropriated to a context other than the one in which it initially evolved, is one mechanism for modular composition. And polymerization is a particularly salient example of modular composition where simpler games evolve to form more complex chains. We also consider how the evolution of new capacities by modular composition may be more efficient than evolving those capacities from basic decisions.

- 1 *Introduction*
 - 2 *The Assembly of Decisions into Games: Cue-Reading, Sensory-Manipulation, and Signalling*
 - 3 *Polymerization*
 - 4 *Template Transfer*
 - 5 *Modular Composition*
 - 6 *Relative Efficiency of Template Transfer*
 - 7 *Discussion*
-

1 Introduction

Evolutionary game theorists analyse equilibria and dynamics of games that are simple enough to be amenable to analysis. The question of the origin of the games studied naturally arises. Social interactions seldom come as neatly structured as the simple games used to model them. So where do the games come from?

On one reading, the question may seem superficial. If this is a problem for the theorist, it seems to be in principle no more of a problem for the game theorist than for any other modeller. A model of an airplane wing in flight or of a bridge with traffic and wind is of necessity a simplification of the complex

real system being modelled. The theorist tries to capture the essence of the phenomena in a form simple enough to be amenable to analysis. This can be done poorly or well—end of discussion.

But there is another version of the question that is deeper. Games model repeated problems of interactive decision. We can ask: How do these games themselves form? How do the bits and pieces of decisions assemble into games? How did individuals, in dealing with the world, come to interact in such a way that the interactions can usefully be characterized as a game with specified players, strategies, information, order of play, and payoffs? And, more generally, how might such games themselves evolve? These are theoretical questions worth investigating.

Our focus here will concern how decisions assemble into games, how games may come to be appropriated to contexts that differ from those in which they are initially played, and, in particular, how simple games assemble into more complex games. There are other aspects of the topic that call for consideration, but we take this to be a suitable place to start.¹

In what follows, we consider a number of types of self-assembly. We concentrate on signalling, although the general perspective applies to other kinds of interaction as well. One might, for example, use the framework discussed here to consider a public goods game evolving to work jointly with a bargaining game. We restrict our attention to signalling games, however, for two reasons. To begin, signalling games are well-studied and the considerations here help to unify recent work on signalling systems in particular. Also, signalling games may evolve to implement simple algorithms that might, by composition, evolve to implement more general forms of computation. We will consider examples of this when we discuss the composition of addition and ordering judgements, and the composition of logical operations.

Throughout, some sort of adaptive dynamics will be fundamental to the analysis. We will consider a small handful of different dynamical processes including simple reinforcement learning, reinforcement with punishment, and reinforcement with invention of new behaviours.²

There will also be a basic structure that characterizes each game. This structure will specify the possible situations that may obtain, the stimuli to which the agents are sensitive, the responses to which they are capable, and the rewards engendered by the combination of actions in the situation at hand. The dynamics indicates how agents' dispositions to respond to stimuli evolve

¹ Note that decisions, at least in the sense we have in mind here, are just agent actions or behaviours that might be reinforced. As such, they presuppose no deliberation or even the capacity for rational reflection.

² See (Erev and Roth [1998]; Barrett and Zollman [2009]; Skyrms [2010]) for discussions of these and other closely related evolutionary processes. Regarding the adaptive dynamics employed in the models, other things being equal, we use the simplest that readily illustrates the particular phenomena of interest.

over time. The game structure itself may evolve, as it does in the case of reinforcement with invention. And the agent's dispositions to respond may be appropriated to new tasks. These new tasks may involve taking as input new aspects of nature or the actions of agents in other games. In addition to evolving new games, the modular composition of games may, under appropriate circumstances, provide more efficient paths towards the evolution of complex dispositions.

2 The Assembly of Decisions into Games: Cue-Reading, Sensory-Manipulation, and Signalling

Consider the origin of signals by the process of ritualization suggested by Tinbergen ([1952]), Lorenz ([1966]), and Huxley ([1966]). Individuals of type *S* just happen to produce a cue that individuals of type *R* learn to read. Then an adaptive dynamics sets up a positive feedback that ritualizes the cue, so that it is used as a signal.

For instance, an animal may naturally bare its teeth before biting, with feedback amplifying the behaviour so that it becomes a ritualized threat signal. The process may lead to complicated signals in simple animals. Scott *et al.* ([2010]) investigate complex vibratory signals of ownership in the caterpillar *Drepana acurata*. These combine vibrations produced by anal scraping, using specialized organs, with those produced by drumming or scraping their mandibles. Using comparative morphology and comparative territorial behaviour across many species, they conclude that anal scraping is a ritualized form of the vibrations made by ordinary locomotion and the vibrations produced by the mandibles are ritualized fighting behaviour. In this species they are combined to form a warning signal of ownership of territory that will be defended.

Suppose that there is a situation where different potential ritualized signals may be sent, depending on the state of the sender, and different behaviours may be appropriate if the receiver reads those signals well. The situation might involve fighting, mating, or anything else that might influence the actions of the receiver.

A particularly salient example is that of the predator-specific alarm calls that are found in many species.³ In the process of becoming ritualized signals, a number of decisions on the part of sender and receiver become bundled together as a game. There is a state of the sender, which nature chooses according to some probability. The sender performs an act that sends a signal, where the signal sent (or better, its probability of being sent) depends on the sender's state. The receiver observes the signal and chooses some act with a

³ See, for example, (Cheney and Seyfarth [1992]; Manser *et al.* [2002]).

probability that depends on the signal received. The act has payoff consequences for both sender and receiver. The individuals involved certainly need not think of this as a game; they may not think at all.

Signalling games may form in a variety of ways. There is evidence that the males of some species may have evolved signal types that exploit pre-existing female sensory biases for the purpose of mating. One well-studied example is the *Physalaemus pustulosus* species group of frogs. The females across the species in the group share preferences for a particular set of call features. One of these, for example, is the presence of low-frequency ‘chucks’. While males in one species may exploit one of these pre-existing preferences, males in a sister species may exploit another. It is argued that the best explanation is that the males in each species have evolved calls that exploit the shared pre-existing preferences of the females and, as it happens, have ended up exploiting different preferences in the context of the different species.⁴

The formation of a signalling game by the exploitation of a pre-existing sensory bias might also be understood as ritualization. The female frogs are attracted by a particular sound, so the male frogs evolve to send mating signals that incorporate that sound. Just as a receiver might evolve to exploit the fixed dispositions of a sender to act in a particular way given the state of nature, here it is the sender who evolves to signal in a way that exploits the receiver’s pre-existing dispositions to act in a particular way given, in this case, a particular type of sound. And a game may be formed by the process of ritualization with positive feedback.

But to speak of ‘positive feedback’ here is to trade specificity for vagueness, as a number of mathematically precise realizations of evolution, imitation, and various types of learning have been rigorously analysed for a wide variety of simple signalling games.⁵ In such games, the positive feedback mechanism is represented by a specific evolutionary or learning dynamics that determines the payoffs for the senders and receivers, and how these payoffs influence their further actions. It is this dynamics that may ultimately forge a ritualized signalling system. When it does, one has an evolutionary model for the formation of that system.

Three simple two-agent games are suggested by such considerations. The first is a cue-reading game, the second a sensory-manipulation game, and the third a signalling game. In a cue-reading game, an agent evolves to take advantage of nature’s fixed dispositions or the fixed dispositions of another agent who comes to play the role of a sender. In a sensory-manipulation game, a sender evolves to exploit the fixed dispositions of an agent who comes to play

⁴ See (Ryan and Rand [1993]). See also (Endler [1993]; Dawkins and Guilford [1996]) for further discussion of the exploitation of pre-existing sensory biases in the evolution of signalling.

⁵ See, for examples, (Hofbauer and Huttegger [2008]; Argiento *et al.* [2009]; Barrett and Zollman [2009]; Skyrms [2010]; Hu *et al.* [2011]; Barrett [2013a]).

the role of a receiver. And each of these games might be understood as simplified versions of a signalling game where a sender co-evolves to respond to nature and a receiver co-evolves to respond to the sender's ritualized signals in a way that is advantageous to both.

In a two-agent cue-reading game, an actor performs some regular fixed action in response to a particular corresponding state of nature. The exploiting agent then observes this action and performs an action of his own. If the exploiter's action is successful, he reinforces his disposition to perform that type of action conditional on the actor's performing the corresponding action. And if it is unsuccessful, he does not reinforce, and may even weaken, this conditional disposition. Here the exploiter may quickly evolve to cue his actions to the behaviour of the actor and hence take advantage of the actor's fixed dispositions to respond to nature.⁶

In a sensory-manipulation game, the sender evolves to take advantage of the fixed dispositions of the receiver in a similar manner. The sender observes the state of nature and sends a signal. The receiver acts in a fixed way on the signal with no choice. If the action is successful, the sender reinforces her dispositions to send that type of signal on the present state of nature. And if it is unsuccessful, she does not reinforce, and may even weaken, this conditional disposition.

In a signalling game, the sender observes the state of nature, then sends a signal. The receiver observes the signal, then performs an action. They may begin by randomly signalling and acting. If the receiver's action is successful, then the sender reinforces the signal sent, conditional on the state of nature that obtained, and the receiver reinforces his action, conditional on the signal he observed. And if the action is unsuccessful, the agents do not reinforce, and may even weaken, these conditional dispositions. In some cases, one can prove that such signalling games converge to signalling systems, and in many cases, the agents evolve successful signalling on simulation.⁷ Here the signals become ritualized through the reinforcement of the sender's dispositions to respond to nature with signals and the receiver's dispositions to respond to signals with actions.

⁶ Where the actor and exploiter have few available types of action, even simple Herrnstein reinforcement learning (Herrnstein [1970]) very quickly leads to successful cue-reading on simulation. Note that the evolutionary task here is typically much easier than the task of coordinating signals and acts in a signalling game inasmuch as the actions of the actor are already fixed in a cue-reading game. The evolution of perception may be modelled by such a game. Indeed, a one-agent cue-reading game, where the agent takes her cues from nature directly, provides a plausible account for how the actor may have evolved her fixed dispositions in response to the various states of nature in the first place. See (Isaac [2011], Section 5.4.1) for a discussion of perception and the information content of perception that fits well with a one-agent cue-reading model.

⁷ See (Skyrms [2010]) for an extended discussion of such games and their properties.

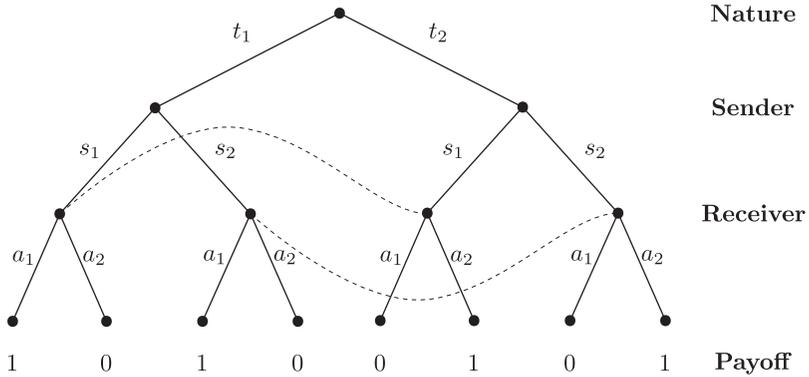


Figure 1. A signalling game.

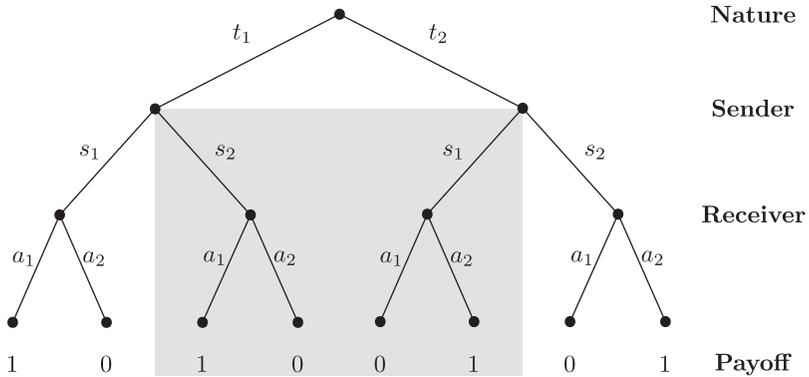


Figure 2. A cue-reading game.

As illustrated in Figure 1, in a Lewis signalling game the first player, the ‘sender’, knows something that the second player, the ‘receiver’, does not know. This is represented by ‘nature’ choosing a state t_1 or t_2 . The sender can only send an arbitrary signal s_1 or s_2 . The receiver observes the signal and takes an action. The interaction is successful, and the payoff to each player is 1, if the receiver chooses a_1 when nature picked t_1 or if the receiver chooses a_2 when nature picked t_2 . Otherwise, the payoff is 0.

A cue-reading game is a Lewis signalling game that has been pruned as indicated by the shaded region of Figure 2. Here the sender has no choice and displays cue s_1 in state t_1 and no cue, which we will indicate by s_2 , in state t_2 . The receiver observes the cue and takes an action.

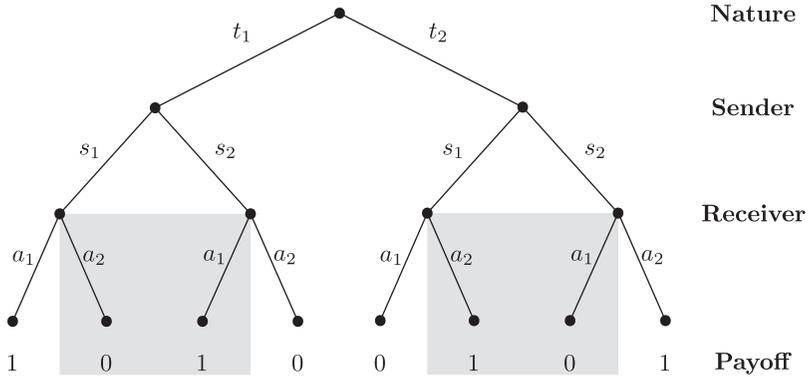


Figure 3. A sensory-manipulation game.

A sensory-manipulation game is a Lewis signalling game that has been pruned as indicated by the shaded region of Figure 3. Here the receiver has no choice and does act a_1 on signal s_1 and act a_2 on signal s_2 .

Just as the cue-reading and sensory-manipulation games can be understood as simplified versions of a Lewis signalling game, they might also be understood as evolutionary steps in the direction of a signalling game as the agents' decisions are ritualized.⁸

Ritualization, then, may act to form cue-reading, sensory-manipulation, or signalling games from decisions. Such games may, then, be extended as other agents exploit established behaviours or form new signalling connections by ritualization. In this sense, one might think of the ritualization of decisions as the glue that binds agents to form simple games from their basic decisions, then increasingly complex games from simple games.

3 Polymerization

The optimal act on receiving a signal may include sending another signal to one or more other agents. Suppose that some individuals learn to do this. Then, in appropriate circumstances, they may spontaneously self-assemble to form chains of senders and receivers to pass information along. For example, consider alarm calls that indicate that a dangerous predator is on the prowl nearby. Individuals who have not seen the predator may repeat the alarm, and pass along the information.

⁸ Cue-reading, sensory-manipulation, and signalling games may differ from each other by degree. A cue-reading game where the sender's initial dispositions are relatively fixed might look increasingly like a signalling game if both the sender and receiver are rewarded for the receiver's actions, further reinforcing the sender's initial dispositional bias.

Consider a simple homogeneous single-species signalling chain. The first member of the chain spots the predator and gives the appropriate alarm call. This individual is in a different game from the other participants—observing the predator and giving the call. The others are all in the same game, observing a call given and giving a call in turn. They have adopted the simple strategy of repeating the call. They have self-assembled into a chain or a more complicated network, whose topology is determined by the contingencies of space and time.

Signalling chains may even cross species lines. The Vervet monkeys studied by Cheney and Seyfarth ([1992]) understand alarm calls of the Superb Starling. Hornbills understand the alarm calls of Diana monkeys (Rainey *et al.* [2004]). Some bird species have learned to understand each others alarm calls (Magrath *et al.* [2009]). Here, at least one species has a more complicated strategy in a more complicated signalling game that takes as inputs both the alarm calls of other species and those of its own, and outputs alarm calls of its own species. Then self-assembly proceeds as before. Simple signalling modules are strung together to form signalling chains and signalling networks.

The structure of the signalling network in these examples is evanescent, depending on who is where when. But there are other examples where the structure of the network is endogenous. Given the costs and benefits, individuals who have learned to react to a signal to their advantage and to ‘pass it along’ may, for example, self-assemble to form rings or stars.⁹

Ritualization of the sort that forms cue-reading, sensory-manipulation, and signalling games in the first place may also do the work of connecting established signalling modules. Suppose that sender S and receiver R may have established a signalling system by ritualization. A new agent, E_1 , who is sensitive to the signalling actions of S may, by reinforcement learning on his own success and failure, learn to read cues and benefit from S 's signals just as S and R do. Then another agent, E_2 , who attends to the actions of E_1 may learn to read cues to his profit. And so on. As new agents are added, the composite game grows in complexity and allows for increasingly subtle connections between agents.¹⁰

The formation of signalling chains has been simulated.¹¹ If one simply places players in a long chain and requires them to learn to signal *de novo*,

⁹ See (Bala and Goyal [2000]) for the description of a ring-forming game and (Huttegger and Skyrms [2013]) for a proof that such networks form under the trial-and-error dynamics of probe and adjust.

¹⁰ Note that what passes as an agent in such games is very simple. It is nothing more than a system that might condition its actions on the actions of other systems. As will be particularly relevant in the next section, an agent may be profitably thought of as a functional unit, among other functional units, in a single organism.

¹¹ See (Skyrms [2009]) for a discussion of how signalling chains may evolve.

it may take a very long time.¹² In some cases, however, the successful evolution of chains may be certain to happen in the limit for chains of arbitrary length.¹³ And if individuals have already learned their strategies in simpler contexts and self-assemble, the evolution of a chain can be quite rapid (Skyrms [2009]).

4 Template Transfer

Template transfer explains how a game that evolved in one context might come to be used successfully in a new context.

When in equilibrium, the agents playing a signalling game have stable dispositions. These stable dispositions might be understood as implementing a rule that takes whatever stimuli to which the senders are sensitive as inputs, and which outputs the actions of the receiver. Template transfer occurs in a signalling game when such an evolved rule is appropriated to a context different from that in which it initially evolved. It involves the evolution of an analogy between the stimuli of the old evolutionary game and a new set of stimuli that characterizes a new game. In many cases the appropriation of an old rule to a new context may be significantly more efficient than evolving a new rule from scratch.

An example of template transfer is illustrated in the transitive-inference behaviour of Pinyon Jays *Gymnorhinus cyanocephalus* and Scrub Jays *Aphelocoma californica*.¹⁴

In an experiment reported by Bond *et al.* ([2003]), seven stimulus colours were arranged in a random linear order that was fixed for each bird. The birds were then presented with two keys, each illuminated with a different colour. If a bird pecked the key illuminated with the higher-ranked colour, then it was rewarded.

The birds were initially presented with only adjacent colour pairs: red and green, green and blue, or cyan and orange, and so on, and the position of the higher-ranked stimulus randomized between left and right keys on each trial. New colour pairs were gradually added as the birds exhibited success in correctly selecting higher-ranked colours. Each of the birds was eventually required to track all six adjacent colour pairs. The Pinyon Jays learned to choose the higher-ranked colour better than 0.85 of the time. The Scrub Jays learned more slowly, but eventually reached a similar level of accuracy.

¹² Even a chain of just three players where the new player is placed between the other two, takes a very long time to evolve (Skyrms [2009]).

¹³ For a discussion of this see Jonathan Kariv's ([2014]) PhD thesis on binary signalling chains.

¹⁴ This section reports selected results from (Barrett [2013b], [2014a], [2014b]). See those papers for further details regarding the model described in this section. See (Barrett [2013b]) as well for an alternative model.

The birds were then presented with non-adjacent colour pairs to determine whether the birds would order the non-adjacent colour pairs based on what they had learned from their experience with just the adjacent colour pairs. Both species immediately exhibited a high level of accuracy in the trials involving the non-adjacent colours, where the experimenters understood accuracy as judgements that matched the colour order that the experiments had initially assigned to the colours. The Pinyon Jays chose the higher-ranked colour in the full ordering with an accuracy of 0.86 for non-adjacent pairs, and the Scrub Jays with an accuracy of 0.77.

Interestingly, the two species exhibited different types of error and different latencies depending on the particular elements in the full colour ordering that were presented. This suggests that the two species employ different mechanisms in making their ordering judgements for non-adjacent pairs. Nevertheless, both species did well judging non-adjacent colour pairs after being trained on just adjacent colour pairs.

The experimenters concluded that the birds were making transitive inferences based on their experience in the first part of the experiment. But the birds were doing more than that. In particular, since an ordering on adjacent colour pairs does not determine an ordering on non-adjacent pairs, the birds were imposing a linear ordering on their learned ordering of adjacent colours. It is only by appropriating a pre-existing linear template that the birds could get from their experience with adjacent colour pairs to judgements that immediately agreed with the experimenters predetermined full linear order. Indeed, that the experimenters themselves took the birds' judgements on non-adjacent colour pairs to be correct and simply inferential suggests that the experimenters were also appropriating a pre-existing linear template to their understanding of the birds' experience.

The evolution of an ordering rule and the subsequent appropriation of that rule to order a new type of stimuli can be modelled by a simple signalling game and its appropriation to a new context.

Consider a signalling game with two senders, A and B , and a receiver R . The senders and receiver may be understood as functional elements of a single bird. Nature chooses at random and without bias two stimuli from a set of seven, and one is presented to each sender, a to A and b to B . The stimuli, which might be represented by the natural numbers 1 through 7, are ordered, with the order represented in the payoffs of the game. The two senders react to the stimuli by sending signals to the receiver, who performs one of three types of act: (0) $a > b$, (1) $a < b$, or (2) $a = b$. The receiver's act will count as successful if and only if it correctly represents the predetermined order of the stimuli to the two agents A and B . The agents evolve an ordering template here if they co-evolve signals that represent the possible stimuli and a dispositional rule for

linearly ordering them that typically produces successful actions by the receiver.¹⁵

Suppose that the agents learn by reinforcement with invention.¹⁶ On this dynamics, one might imagine that each sender has an urn corresponding to each possible stimulus and that each urn begins with just a single black ball. When presented with a particular stimulus, each sender draws a ball at random from the corresponding urn. If the ball is black, a new signal type is invented and sent to the receiver; otherwise, a signal of the type of the drawn ball is sent. The receiver also has an urn corresponding to each ordered pair of signals he might receive. And if he gets a new signal type, he introduces corresponding urns to represent the new possible pairs. Each of the receiver's urns begins with a single ball of his three act types: $a > b$, $a < b$, and $a = b$. If successful, the ball drawn from each urn is returned and a new ball of the signal or act type used in that play of the game is added to the urn; otherwise, the ball drawn from each urn is just returned. Finally, newly invented signal types are only kept and reinforced if they lead to a successful act the first time they are used. Thus the game itself evolves.

On simulation, the senders begin by quickly inventing an assortment of new signals that they initially send more or less randomly. Consequently, the receiver initially acts pretty much at random. 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.

The composite system might be thought of as having evolved to implement a dispositional rule that takes naturally ordered stimuli as input, represents the stimuli as signals, then outputs an act that reliably indicates the natural order of stimuli. Once evolved, this dispositional rule might be appropriated as a template to represent and judge the natural order of a new type of stimuli.

Such a template might be fit to a new context by coordinating the new stimuli with the old inputs to the dispositional rule. The association of the new stimuli with the old inputs to the dispositional rule might be thought of as implementing an analogy between the new and old stimuli. When such an analogy evolves, the old dispositional rule evolves to treat the new stimuli similarly to how it treated the old stimuli that were involved in forging the old dispositional rule. This sort of template transfer might be evolutionarily favoured when the process that coordinates the new stimuli to the old inputs is more efficient than evolving a new rule for the new context.

¹⁵ Compare to Ariel Rubenstein's discussion of linear order in (Rubenstein [2000], Chapter 2).

¹⁶ See (Argiento [2009]; Skyrms [2010]; Alexander *et al.* [2012]) for descriptions of this dynamics and its properties. We consider this adaptive dynamics here as it shows how the agents might both invent a representation for the stimuli and develop successful ordering dispositions with modest evolutionary resources. This invented representation is then appropriated to a new task.

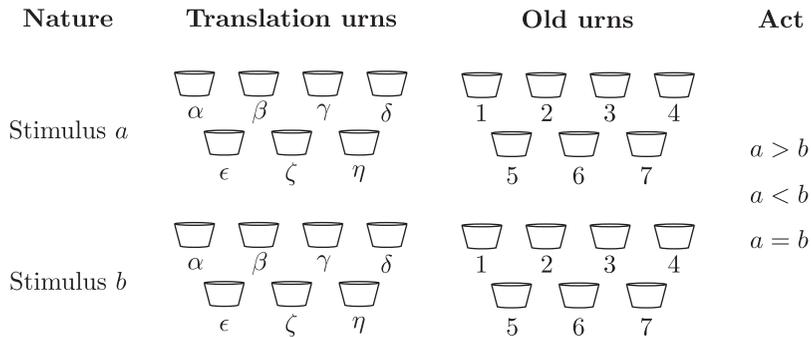


Figure 4. A simple example of template transfer.

We investigate this possibility by considering a simple model. Consider a new ordered set of stimuli represented by the first seven letters of the Greek alphabet α – η . Suppose that the agents in the last game have already evolved a dispositional rule for representing and ordering the old stimuli, 1–7, but that they must learn to represent and linearly order the new Greek stimuli to be successful. The old evolved dispositional rule might be used as a template for ordering the new stimuli if the agents can evolve to associate the new stimuli to the old inputs in a one-to-one way.¹⁷

Suppose that nature chooses at random and without bias two stimuli from the set α – η , and one is presented to each sender, *A* and *B*. Each sender has an urn corresponding to each of the new stimuli types (Figure 4). Each of these urns initially contains one ball corresponding to each of the old sender urns, 1–7. When presented with a new stimulus, each sender draws a ball at random from one of her new Greek urns, then draws a ball from the old Arabic urn indicated on the ball from the first draw, then sends the signal indicated on the second draw to the receiver. Since the agents have already evolved dispositions that order the old stimuli (and hence the old Arabic urns), the agents just need to learn to associate the new Greek urns with the old Arabic urns in a way that respects the new linear order to be successful on the new stimuli.

We will suppose that the agents learn in this new game by simple reinforcement with punishment, where success is determined by whether the results of the receiver's judgements match the natural order of the new stimuli.¹⁸ More

¹⁷ If there were fewer new stimuli than old stimuli, the agents might be successful even if their association of the new stimuli to the old inputs was not one-to-one, but still respected the natural order of the new stimuli.

¹⁸ Our strategy regarding the adaptive dynamics is to use the simplest sort for the explanation at hand. Using reinforcement with invention to evolve the ordering dispositions in the basic game explains how the birds might invent an internal representation and the associated dispositions. Reinforcement with punishment explains how this invented system might be transferred to a new context on a simple dynamics.

specifically, if, on a particular play of the game, the agents choose balls from Greek urns that indicate Arabic urns that produce signals that lead the receiver to correctly order the new stimuli, then each sender returns her ball to the Greek urn from which it was drawn and adds a copy of the same ball type; otherwise, she 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.¹⁹

On simulation, the agents typically (0.995 on 1000 runs) evolve to successfully match the new Greek stimuli to the corresponding old Arabic ordering system with an accuracy better than 0.8 with 10^5 plays per run. Here the receiver's dispositions are already well tuned to making successful linear ordering judgements on the old stimuli. The senders, then, just need to evolve a successful analogy between the old and new stimuli.

With respect to modelling the behaviour of the jays, when the old dispositional rule evolves in the context of ordering Arabic inputs and the new association urns are trained on just adjacent Greek inputs, the new composite system evolves to judge the full order of Greek stimuli with an accuracy better than 0.80. Here the composite system sometimes evolves to match Greek to Arabic stimuli in a one-to-one manner that respects the full Arabic ordering, but more typically the Arabic ordering rule is used to group the Greek stimuli into two or three linearly ordered segments. Within each segment, the composite system does very well in ordering both adjacent and non-adjacent Greek stimuli. Between segments, the judgements are less reliable, but in aggregate, the composite system does about as well as the judgements of the jays when they apply a prior linear-order template to their experience of only adjacent colour pairs.²⁰

5 Modular Composition

Modular composition occurs when a game comes to accept the play of another game as input, thus forming a composite game. Here the stimuli to which the agents in one game are sensitive are the actions of agents playing other games. Polymerization is a special case of modular composition. More generally, modular composition of simple games may lead to heterogeneous networks of arbitrary topology.

¹⁹ Note that the contents of the old arabic-ordering urns do not change on plays of this game, and, hence, the template is assumed to be fixed. One could allow the template to itself evolve as it is transferred. This might better tune the template to the new context, but that would also involve a different game than the one discussed here.

²⁰ See (Barrett [2014b], Section 6) for further details regarding how an old ordering template evolves to order a new type of stimuli on the incomplete evidence of only adjacent pairs of the new stimuli.

There are two aspects to modular composition. The first concerns how the inputs to the component modules are determined. The second concerns how modules come to interpret the actions of other modules to allow for successful coordinated action.

Concerning the first, the natural saliences for an agent are determined by the perceptual apparatus of the agent. For the purposes at hand, we will simply stipulate what each agent responds to, then consider how they might evolve to use these inputs and other evolved dispositions for successful action. Of course, an agent's natural saliences might themselves evolve on feedback from her success and failure in action.

Concerning the second, games may come to interpret the actions resulting from the play of other games by the same reinforcement mechanism that allows for template transfer. In particular, the glue that binds a simple game to other games is an evolved analogy between how the game treated its old inputs and how it treats new inputs resulting from the play of the other games. Here is a rich experimental example, accompanied by a simple computational model.

In a recent series of experiments, Livingstone *et al.* ([2014]) trained Rhesus macaque monkeys to associate the number of dots on either side of a touch screen with the number of food drops of reward the monkey would get if it selected that side of the screen. They then trained the monkeys to associate symbols with the number of food drops, to add the symbolically represented reward magnitudes, and to transfer their learned arithmetic competence to a new set of symbols. We will briefly discuss the experiments, then model the behaviour of the monkeys by the modular composition of two games. One game will involve ordering judgements, the other addition.

The first experiment was a dots comparison task. Here the monkeys were presented with two circles on either side of the touch screen each containing randomly placed dots of various sizes and colours. When a monkey touched a side of the screen, it was rewarded with the number of liquid drops corresponding to the number of dots on the side it chose. The device that fed the monkeys distinctly beeped once as each drop was dispensed. The monkeys learned to choose the option with the greater number of dots with an accuracy between 80% and 90%.

They were then trained on the first symbol comparison task. Here the experimenters associated a particular symbol with each cardinality from 0 to 25. The symbols were neither cognates for the numbers nor did they exhibit any special pattern. The monkeys were then presented with a symbol on each side of the touch screen and were rewarded with a number of drops, each accompanied by a beep corresponding to the predetermined value of the symbol they selected. The monkeys learned to choose the option with the greater number

of dots with a high degree of accuracy, particularly when there was a significant difference between the value of the two symbols.

In the dots addition task, the monkeys were presented with one set of dots on one side of the screen and two sets of dots on the other side of the screen, where each set was specified by a circle around the dots. They were then rewarded in drops, accompanied by beeps, for the total number of dots on the side of the screen that they selected. Here the monkeys learned to reliably pick the larger of the two options when the sum of cardinalities of the sets on one side of the screen was significantly different from the cardinality of the set on the other. When the two options were close in value, their behaviour approached chance.

For the first symbol addition task, the monkeys were presented with two of the symbols they had learned in the first symbol comparison task in an oval on one side of the screen and a single symbol on the other side of the screen. The monkeys were then rewarded in drops, accompanied by beeps, equal to the sum of the symbols in the oval or the value of the single symbol, respectively. Significantly, the monkeys initially acted as if they were comparing the larger of the two addends against the symbol on the other side of the screen. But they eventually learned to combine the magnitudes represented by the two addends and reliably pick the larger of the two options when their sum was significantly different from the cardinality represented on the other side of the screen.

Once the monkeys were successful at this task, a new set of symbols was introduced in the second symbol comparison task. They successfully learned this new set of symbols by comparison, just as they had learned the first set of symbols in the first symbol comparison task.

Finally, in order to determine whether the monkeys could transfer their arithmetic competence with the old symbol set to the new symbol set, they were presented with a second symbol addition task where they were required to add the new symbols just as in the first symbol addition task. The experimenters argued that if the monkeys could learn to do so more efficiently than in the first symbol addition task, this would indicate that they were able to transfer their arithmetic competence to a new system of representation, which might then be taken as evidence that they were using the symbols as representations to carry out computations and not just memorizing particular symbol combinations.

In the first symbol addition task, the monkeys reached a stable success rate after about fifty days. But when presented with the second symbol addition task, the monkeys immediately chose the larger side more often than they did at a comparable time during the first symbol addition task, and their performance reached a stable success rate in just ten days. The experimenters interpreted this as evidence that the monkeys were able to transfer their previously learned competence at symbolic arithmetic to a new context and, hence, were

carrying out computations on the cardinalities represented by the symbols and not just memorizing symbol combinations.

While it is unclear precisely what the monkeys were doing cognitively in each task, one might model the evolution of such arithmetic competences using signalling games, template transfer, and modular composition.

The dot comparison and the first and second symbol comparison tasks involve the monkeys learning orderings of collections of dots and symbols just as the jays learned orderings of colours. This can be modelling by a two-sender signalling game similar to the one discussed in connection with the jays. Here, each of the two senders in the model, again thought of as different functional elements of a single animal agent, would have access to one side of the screen.

Template transfer illustrates how an ordering rule that evolved in one context might come to be used successfully in another context, as in the second symbol addition task where the monkeys learn to use the second symbol set in a way analogous to the first symbol set. Indeed, by means of template transfer, the monkeys might learn the second symbol addition task more efficiently than they learned the first symbol addition task, without learning to carry out arithmetic computations. More specifically, they might learn successful combinations of symbols from the first symbol set, then evolve an analogy between the old symbol set and the new symbol set by template transfer.

If so, this poses a challenge to the experimenters' conclusion that the second symbol addition task involves computation. In the model just described, the competence of the monkeys in the second symbol addition task is achieved by memorization and template transfer, not computation. Further, template transfer here also explains why the monkeys learned the second symbol addition task faster than the first. On the other hand, transferring a learned rule of association from the context in which it was learned to a new context by analogy might count as evolving a very basic type of concept. Here these concepts might represent cardinalities and how they might be ordered.²¹

There are several ways that the monkeys might evolve to order dot-sums and symbol-sums in the first place. They might evolve the ability to add the cardinalities and order them all at once. Or they might learn how to add and how to order separately, then learn how to compose these skills to carry out the addition-ordering tasks set for the monkeys here.

In this second case, the dot addition and the first symbol addition tasks might be modelling as the modular composition of two, more basic, evolutionary games. More specifically, one might consider the composition of a two-sender ordering game very much like the one discussed in the last section

²¹ See (Barrett [2014b]) for a discussion of how the evolved appropriation of a rule might be understood as the evolution of increasingly general concepts in the context of the ordering game discussed above.

and a two-sender addition game that evolves to compute the sums of the cardinalities presented to each sender.

We will consider the two-sender addition game first, then consider how that game might evolve to communicate successfully with the two-sender ordering game.

In the two-sender addition game, each sender is presented with a random set of zero to five dots (or, equivalently, one of six symbols), where each cardinality has equal probability. The senders then each send a signal to the receiver who performs an action that depends on the signals. If that action corresponds to the sum of cardinalities of the dots presented to the senders, then that particular play was successful and the three agents, who can together be thought of as functional parts of single composite agent, reinforce the conditional actions just taken.

We will suppose, as we did with the evolution of the ordering rule in the last section, that the agents learn by reinforcement with invention.²² Each sender has six urns, one for each cardinality 0–5. Each urn initially contains just a single black ball. While there is always precisely one black ball in each of the senders' urns, other balls, labelled with terms corresponding to potential signals, are added to the urns on repeated plays of the game. When presented with a set of dots, each sender draws a ball at random from the corresponding urn. If the ball is black, the sender invents a new term at random and sends that term to the receiver; alternatively, if the ball is labelled with a term, the sender simply sends that term to the receiver. On each play of the game, then, the receiver gets one signal from each sender.

The receiver has an urn labelled by each pair of terms that he has received from the senders so far, and he constructs further new urns as needed when the senders invent and send new terms. Each of the receiver's urns begins with 11 balls, one ball corresponding to each possible sum action from 0 to 10. When the receiver gets the pair of signals from the sender, he draws a ball at random from the urn corresponding to that particular pair of signals, then carries out the corresponding act.

If the act matches the sum of the cardinalities of dots and neither sender drew a black ball, then each agent returns the ball he drew to its urn and adds a copy of that ball. If a sender drew a black ball and the receiver's action was successful, then the sender adds one ball of the new signal type he invented to each of his urns, and the receiver adds urns corresponding to the new types of signals he may receive. If the receiver's action was not successful, the agents just return the balls they drew to the urns. Note that once a new term is invented, the agents learn by simple reinforcement without punishment.

²² Just as in the case of the order rule in the last section, using reinforcement with invention as the adaptive dynamics here explains how the an appropriate representation of states and actions might be invented.

On simulation, since the senders do not initially have much in their urns, they draw the black ball often and, hence, begin by inventing new terms at a relatively high rate. While they initially send the newly invented terms at random, and the receiver initially acts randomly when he receives a new pair of terms, as the agents update their first-order dispositions by reinforcement, they typically evolve a set of systematically interrelated dispositions where the receiver's act corresponds to the sum of the cardinalities presented to each of the senders.²³ Here the senders' terms have evolved to represent cardinalities and the receiver's dispositions have evolved to add those cardinalities.

As with the birds, the agents in the game might be taken to represent various functional parts of a single monkey. On this interpretation, the model illustrates how a monkey might co-evolve an internal representation of cardinalities and the ability to reliably calculate sums on the basis of this representation.²⁴

Given that the monkeys possess such addition dispositions—which, as we have just seen, might themselves evolve in the context of a signalling game—one might model the behaviour of the monkeys in the dot addition and the first symbol addition tasks as the modular composition of a two-sender ordering game and a two-sender addition game. Here the addition game takes its two inputs from the two sets of dots on one side of the screen, then the ordering game takes one input from the output action of the addition game and the other input from the dots on the other side of the screen. If these games can evolve to communicate successfully, the result will be a composite game where the monkeys sum the cardinalities on one side of the screen, then compare this sum against the cardinality on the other side of the screen and select the larger result.

Suppose that the addition game has evolved to compute the sum of inputs a and b from the left side of the screen and produce an act corresponding to the sum. And suppose that the comparison game has evolved to choose the greater of two cardinalities, l and r , and is sensitive to the act $a + b$ of the addition game and the cardinality of the right-hand side of the screen, r . The comparison game has evolved to add cardinalities of sets of dots, so, at least initially, it does not know what to do with the act $a + b$ of the addition game. In order to work together, an association must evolve between the two modules. This can occur as a form of template transfer.

²³ In particular, on 1000 runs of 8.0×10^6 plays each, the cumulative success rate is greater than 0.95 in 0.67 of the runs, greater than 0.90 in 0.83 of the runs, and greater than 0.85 in 0.95 of the runs.

²⁴ See (Barrett [2013a]) for a model where the agents evolve the ability to add increasingly large cardinalities. The sense in which this counts as an arithmetic computation is relatively weak, but insofar as the agents can learn to transfer this rule to other contexts, it might, again, be thought to represent a potentially general rule and hence an arithmetic computation.

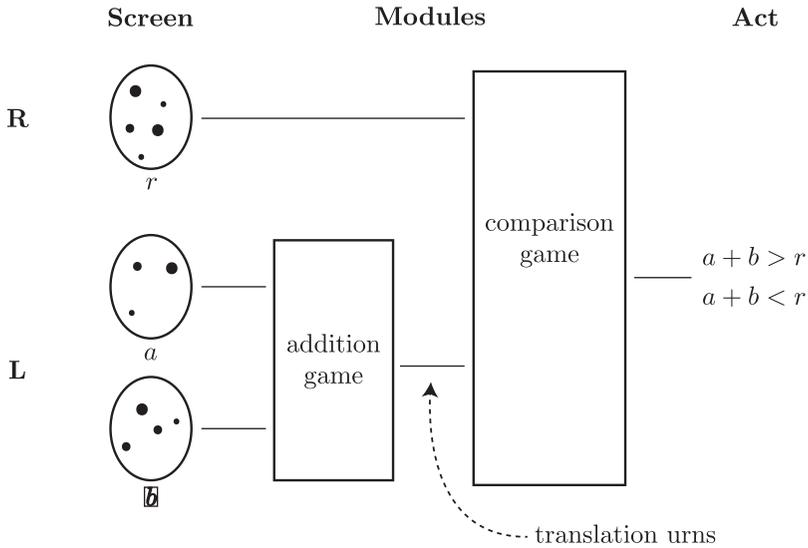


Figure 5. Modular composition of addition and ordering.

When a set of dots is presented to the R input, the sender in the comparison game sends whatever signal it has evolved to send. But when an act from the addition game is presented to the L input of the comparison game, it is matched up, at least initially, with a random signal.

Just as in the bird-ordering template transfer model, we will consider how this analogy might evolve by reinforcement with punishment. Consider eleven translation urns labelled 0 to 10, corresponding to each of the possible act types resulting from the addition game. Each urn initially has one ball of each signal type used in the comparison game labelled 0 to 10. When an action is presented to the L input of the comparison game, a ball is drawn from the corresponding urn and the signal corresponding to that ball is sent to the receiver of the comparison game (Figure 5). If the comparison game judges the relative cardinalities of $a + b$ and r correctly, where either $a + b > r$ or $a + b < r$ in each case, then the translation ball drawn is replaced in the urn from which it was drawn and a copy is added; otherwise, the ball drawn is discarded unless it is the last ball of its type in that urn.

On simulation, the addition and comparison games are found to be successfully coordinated by reinforcement learning with punishment. With 1000 runs of 10^5 plays each, the composite system evolved to correctly order $a + b$ and r with an accuracy typically about 0.94 and always better than 0.90. And when it gets the order wrong, the mean difference between $a + b$ and r was 1.52. This agrees well with the monkeys' degree of success, and with their being

more likely to make a mistake when the cardinalities being compared are close.²⁵

Part of the work is done by how the two modules are connected. Here we simply stipulated in the specification of the composite game the stimuli to which the inputs to the addition and ordering rules were sensitive, without saying how the modules might have come to be sensitive to those particular types of stimuli. There are, however, considerations that may serve to limit the options.

In the case of chains of alarm calls, the question of what each agent attends to was unproblematic. Alarm calls have been learned to deserve attention in simple contexts, and one might expect this salience to transfer to more complicated chains or networks of games. Further, the same features that make for a good signal—chiefly, that it is easily noticed given the *de facto* perceptual capacities of the agents—may also make it salient to the agents responsible for the input to other modules. More generally, the actions resulting from evolved modules are actions that have mattered to payoffs. Insofar as agents attend to payoffs, then, such actions might be similarly salient in similar contexts.

In short, whether it is more efficient to evolve a new capacity from scratch or by modular composition will depend to a large degree on precisely what is salient to the agents and how.

6 Relative Efficiency of Template Transfer

The modules in the bird transitive inference and monkey addition examples evolved on reinforcement with invention, and template transfer evolved on reinforcement with punishment. In order to compare the relative efficiency of template transfer as directly as possible, we will consider a very simple logical game that evolves on reinforcement learning, then compare this to template transfer on precisely the same learning dynamics.

The NAND game is a two-sender, one-receiver signalling game where the agents learn by simple reinforcement. Nature randomly and independently picks state 0 or 1 for each sender, A and B . Let S_A and S_B be these states. Each sender has one urn for each of the two possible states, and each urn starts with one ball indicating signal g and one indicating signal r . The receiver, R , has one urn for each pair of signals he might receive from the two senders: $[g_A, g_B]$, $[g_A, r_B]$, $[r_A, g_B]$, and $[r_A, r_B]$, and each of these urns starts with one ball labelled '0' and one ball labelled '1'. A play of the game is successful if and

²⁵ Indeed, one should expect the latter phenomenon when a skill evolves by reinforcement on ordering judgements since the composite agent need not get the order exactly right to enjoy a high degree of success. One might, consequently, expect the monkeys to exhibit more precise addition behaviour in an experiment where the success of their action on each play depends on getting the sum exactly right, rather than just getting an order judgement right.

only if R performs the act corresponding to $S_A \text{ NAND } S_B$ —that is, if R does act 0 when both S_A and S_B are 1, and does act 1 otherwise. If a play is successful, each agent returns the ball she drew to the urn from which it was drawn and adds another of the same type; otherwise, each agent simply returns the ball she drew.

The NAND game typically co-evolves a basic signalling language and the operator NAND. On 1000 runs with 10^6 plays/run, 0.71 of the runs exhibit a cumulative success rate of better than 0.80, 0.62 of the runs better than 0.90, and 0.50 of the runs better than 0.95. The game sometimes does not evolve NAND. When it fails to do so, every pair of states produces act 1 and, consequently, the game is successful on about 0.75 of the plays.²⁶

Once NAND has evolved, it might be appropriated to a new context by template transfer. For the purposes of direct comparison, we will also consider template transfer in the context of simple reinforcement learning.

Suppose that nature presents each sender with random state $0'$ or $1'$. Each sender has two translation urns, one labelled $0'$ and one labelled $1'$, that each start with one ball labelled 0 and one labelled 1. Each sender draws from the translation urn corresponding to the current state presented to her, draws from the old urn indicated on the ball drawn from the new translation urn, then sends the old signal indicated by that draw to the receiver who then acts with whatever dispositions he acquired when NAND evolved in the old context. A play of this game is successful if and only if the receiver preforms the act corresponding to $S_{A'} \text{ NAND } S_{B'}$. If the play is successful, each of the senders returns the ball she drew to the translation urn from which it was drawn and adds another ball of that type; otherwise, each sender simply returns the ball she drew. The receiver does not change his dispositions.

This template transfer game typically appropriates NAND to the new context an order of magnitude faster than the initial evolution of the logical operator on the same dynamics. It does so by evolving a map from new inputs to old signals that exploits the fact that the receiver's dispositions have already evolved to calculate NAND on those signals. On 1000 runs with 10^5 plays per run, 0.78 of the runs exhibit a cumulative success rate of better than 0.80, 0.61 of the runs better than 0.90, and 0.50 of the runs better than 0.95, which is about the same level of success that the original NAND game achieves on 10^6 plays per run.²⁷

Template transfer also allows for a more subtle sort of efficiency. Once NAND has evolved, it may be appropriated to a new context by template

²⁶ This happens on about 0.23 of the runs with 10^7 plays per run.

²⁷ For further comparison, on 10^5 plays per run in the original NAND game, 0.60 of the runs exhibit a cumulative success rate of better than 0.80, 0.48 of the runs better than 0.90, and 0.34 of the runs better than 0.95. Like the NAND game, the template transfer game also fails to evolve NAND for the new inputs on about 0.23 of the runs with 10^7 plays per run.

transfer to play the role of a different logical operator more efficiently than that operator might evolve on its own. We will consider one example of this.

The operation OR typically evolves on simple reinforcement learning just as in the NAND game. And, perhaps unsurprising given the similarity in their truth tables, its evolution exhibits about the same efficiency on simulation.

Consider the template transfer of NAND into a context where a play is successful if and only if the receiver preforms the act corresponding to S_A' or S_B' , and suppose that the senders update their translation urns by simple reinforcement. On simulation, the composite system evolves the logical operation OR from NAND an order of magnitude faster than OR evolves on the same dynamics. Template transfer accomplishes this by mapping new inputs to old signals in such a way that the receiver produces his 0 output in state (0, 0) instead of (1, 1). This works because both logical operations produce 0 on exactly one pair of inputs.²⁸

7 Discussion

The ritualization of decisions explains how cue-reading, sensory-manipulation, and signalling games may initially evolve and how new games may evolve from old by polymerization, template transfer, and modular composition. These processes are interrelated.

The ritualization of decisions explains how cue-reading, sensory-manipulation, and signalling games may get started. An agent may condition his actions on some aspect of nature, then learn from the consequences. An agent may condition his actions on some aspect of nature and another condition her actions on the actions of the first agent, then each learn from the consequences. An agent may learn to manipulate the actions of another agent by taking advantage of the fixed dispositions of the second agent. Or agents may co-evolve dispositions that coordinate their actions to their mutual benefit. In each of these cases, the decisions that result from the *de facto* dispositions of the agents form a game by ritualization. Such games may eventually lead to stable, successful dispositions that implement a rule. And such modular rules may, in turn, compose with each other by further ritualization of decisions.

Polymerization is a simple case of modular composition. Chains of alarm calls provide a compelling example of the sort of natural saliences that may serve to connect modules that initially evolved independently. Given appropriate saliences, ritualization of the sort that leads to

²⁸ This works the same way for the other two binary logical operators that produce one false. More generally, if one had five binary logical operators, one for each possible number of false outputs, one could get the other eleven by template transfer. And each would evolve an order of magnitude faster this way than on its own.

cue-reading, sensory-manipulation, and signalling explains how such modules might compose to allow for successful coordinated action. The process of template transfer, in the examples described here, is a special type of cue-reading that is particularly well suited to modular composition. And, while the examples here involve a type of cue-reading, template transfer might also compose modules by the ritualization of decisions between agents, akin to those that form sensory-manipulation or signalling games.

Template transfer explains how a game or rule may come to function in a context different from that in which it initially evolved. The appropriation of the module to a new context involves the evolution of an analogy between the old inputs to the game and the new inputs. The new inputs might come from nature or from the actions of other agents. In the later case, a module may learn how to interpret the actions of other modules as inputs to produce successful actions. Such functional composition between modules may be more efficient than evolving the new capacity from scratch.

The evolution of strategies in a given game is a vibrant area of ongoing research. But the question of the evolution of games themselves is important and deserves to be explored. Here we have taken a few initial steps.

Acknowledgements

We would like to thank Louis Narens, Simon Huttegger, Kevin Zollman, Kai Wehmeier, Alistair Isaac, Cailin O'Connor, and the two anonymous reviewers for helpful comments. We would also like to thank Thomas Barrett for producing the figures.

Jeffrey A. Barrett
 UC Irvine
 Irvine, CA 92697, USA
 j.barrett@uci.edu

Brian Skyrms
 UC Irvine
 Irvine, CA 92697, USA
 bskyrms@uci.edu

References

- Alexander, J. M., Skyrms, B. and Zabell, S. L. [2012]: 'Inventing New Signals', *Dynamic Games and Applications*, **2**, pp. 129–45.
- Argiento, R., Pemantle, R., Skyrms, B. and Volkov, S. [2009]: 'Learning to Signal: Analysis of a Micro-level Reinforcement Model', *Stochastic Processes and Their Applications*, **119**, pp. 373–90.

- Bala, V. and Goyal, S. [2000]: 'A Noncooperative Model of Network Formation', *Econometrica*, **68**, pp. 1181–229.
- Barrett, J. A. [2007]: 'Dynamic Partitioning and the Conventionality of Kinds', *Philosophy of Science*, **74**, pp. 527–46.
- Barrett, J. A. [2013a]: 'On the Coevolution of Basic Arithmetic Language and Knowledge', *Erkenntnis*, **78**, pp. 1025–36.
- Barrett, J. A. [2013b]: 'The Evolution of Simple Rule-Following', *Biological Theory*, **8**, pp. 142–50.
- Barrett, J. A. [2014a]: 'The Evolution, Appropriation, and Composition of Rules', *Synthese*, doi: 10.1007/s11229-014-0421-6.
- Barrett, J. A. [2014b]: 'Rule-Following and the Evolution of Basic Concepts', *Philosophy of Science*, **81**, pp. 829–39.
- Barrett, J. A. and Zollman, K. [2009]: 'The Role of Forgetting in the Evolution and Learning of Language', *Journal of Experimental and Theoretical Artificial Intelligence*, **21**, pp. 293–309.
- Bond, A. B., Kamil, A. C. and Balda, R. P. [2003]: 'Social Complexity and Transitive Inference in Corvids', *Animal Behaviour*, **65**, pp. 479–87.
- Cheney, D. and Seyfarth, R. [1992]: *How Monkeys See the World*, Chicago, IL: University of Chicago Press.
- Dawkins, M. S. and Guilford, T. [1996]: 'Sensory Bias and the Adaptiveness of Female Choice', *The American Naturalist*, **148**, pp. 937–42.
- Endler, J. A. [1993]: 'Some General Comments on the Evolution and Design of Animal Communication Systems', *Philosophical Transactions of the Royal Society of London*, **B340**, pp. 215–25.
- Erev, I. and Roth, A. E. [1998]: 'Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria', *American Economic Review*, **88**, pp. 848–81.
- Herrnstein, R. J. [1970]: 'On the Law of Effect', *Journal of the Experimental Analysis of Behavior*, **13**, pp. 243–66.
- Hofbauer, J. and Huttegger, S. [2008]: 'Feasibility of Communication in Binary Signaling Games', *Journal of Theoretical Biology*, **254**, pp. 843–9.
- Hu, Y., Skyrms, B. and Tarrès, P. [2011]: 'Reinforcement Learning in a Signaling Game', <<http://arxiv.org/abs/1103.5818>>.
- Huttegger, S. [2007]: 'Evolution and the Explanation of Meaning', *Philosophy of Science*, **74**, pp. 1–27.
- Huttegger, S. and Skyrms, B. [2013]: 'Emergence of a Signaling Network with Probe and Adjust', in B. Calcott, R. Joyce and K. Sterelny (eds), *Cooperation and Its Evolution*, Cambridge, MA: MIT Press, pp. 265–74.
- Huttegger, S., Skyrms, B., Tarrès, P. and Wagner, E. [2014]: 'Some Dynamics of Signaling Games', *Proceedings of the National Academy of Sciences*, **111**, pp. 10873–80.
- Huxley, J. [1966]: 'A Discussion on Ritualization of Behaviour in Animals and Man: Introduction', *Philosophical Transactions of the Royal Society of London*, **251**, pp. 249–71.

- Isaac, A. [2011]: *The Informational Content of Perceptual Experience*, PhD thesis, Stanford University.
- Kariv, J. [2014]: *Broken Telephone: An Analysis of a Reinforced Process*, PhD thesis, University of Pennsylvania.
- Lewis, D. [1969]: *Convention*, Cambridge, MA: Harvard University Press.
- Livingstone, M. S., Pettine, W. W., Srihasam, K., Moore, B., Morocz, I. A. and Lee, D. [2014]: 'Symbol Addition by Monkeys Provides Evidence for Normalized Quantity Coding', *Proceedings of the National Academy of Science*, **111**, pp. 6822–7.
- Lorenz, K. [1966]: 'Evolution of Ritualization in the Biological and Cultural Spheres', *Philosophical Transactions of the Royal Society of London B*, **251**, pp. 273–84.
- Magrath, R. B., Pitcher, B. J. and Gardner, J. L. [2009]: 'Recognition of Other Species' Aerial Alarm Calls: Speaking the Same Language or Learning Another?', *Proceedings of the Royal Society B*, **276**, pp. 769–74.
- Manser, M., Seyfarth, R. and Cheney, D. [2002]: 'Suricate Alarm Calls Signal Predator Class and Urgency', *Trends in Cognitive Science*, **6**, pp. 55–7.
- McGregor, P. [2005]: *Animal Communication Networks*, Cambridge: Cambridge University Press.
- Ouattaraa, K., Lemassona, A. and Zuberbühler, K. [2009]: 'Campbell's Monkeys Concatenate Vocalizations into Context-Specific Call Sequences', *Proceedings of the National Academy of Science*, **106**, pp. 22026–31.
- Rainey, H. J., Zuberbühler, K. and Slater, P. J. B. [2004]: 'Hornbills can Distinguish Between Primate Alarm Calls', *Proceedings of the Royal Society B*, **271**, pp. 755–9.
- Roth, A. E. and Erev, I. [1995]: 'Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Immediate Term', *Games and Economic Behavior*, **8**, pp. 164–212.
- Rubenstein, A. [2000]: *Economics and Language*, Cambridge: Cambridge University Press.
- Ryan, M. J. and Rand, P. J. B. [1993]: 'Sexual Selection and Signal Evolution: The Ghosts of Biases Past', *Philosophical Transactions of the Royal Society of London B*, **340**, pp. 187–295.
- Scott, J., Kawahara, A., Skevington, J., Yen, S., Sami, A., Smith, M. and Yack, J. [2010]: 'The Evolutionary Origins of Ritualized Acoustic Signals in Caterpillars', *Nature Communications*, **1**, pp. 1–9.
- Skyrms, B. [2008]: 'Signals', *Philosophy of Science*, **75**, pp. 489–500.
- Skyrms, B. [2009]: 'Evolution of Signaling Systems with Multiple Senders and Receivers', *Philosophical Transactions of the Royal Society B*, **364**, pp. 771–9.
- Skyrms, B. [2010]: *Signals: Evolution, Learning, and Information*, New York: Oxford University Press.
- Tinbergen, N. [1952]: "'Derived" Activities: Their Causation, Biological Significance, Origin, and Emancipation during Evolution', *Quarterly Review of Biology*, **27**, pp. 1–23.