



# Language games and the emergence of discourse

Jeffrey A. Barrett<sup>1</sup> · Jacob VanDrunen<sup>1</sup>

Received: 2 August 2021 / Accepted: 4 March 2022 / Published online: 15 April 2022  
© The Author(s) 2022

## Abstract

Wittgenstein (Philosophical investigations, Basil Blackwell Ltd., 1958) used the notion of a language game to illustrate how language is interwoven with action. Here we consider how successful linguistic discourse of the sort he described might emerge in the context of a self-assembling evolutionary game. More specifically, we consider how discourse and coordinated action might self-assemble in the context of two generalized signaling games. The first game shows how prospective language users might learn to initiate meaningful discourse. The second shows how more subtle varieties of discourse might co-emerge with a meaningful language.

**Keywords** Language games · Signaling games · Self-assembling games · Evolutionary games · The emergence of discourse

## 1 Introduction

Ludwig Wittgenstein was concerned with how meaningful language was interwoven with action.<sup>1</sup> As he put it, in learning a language “children are brought up to perform *these* actions, to use *these* words as they do so, and to react in *this* way to the words of others” (1958, §6).<sup>2</sup> Learning a language involves establishing an association between words and actions. To illustrate how meaningful language is interwoven with action, Wittgenstein described a simple language meant to serve for communication between a builder *A* and her assistant *B*:

*A* is building with building-stones: there are blocks, pillars, slabs and beams.  
*B* has to pass the stones, and that in the order in which *A* needs them. For this

<sup>1</sup> We would like to thank Brian Skyrms and the anonymous reviewers for their helpful comments on earlier versions of the paper.

<sup>2</sup> The Wittgenstein references are to the numbered sections of his *Philosophical Investigations*.

✉ Jeffrey A. Barrett  
j.barrett@uci.edu  
Jacob VanDrunen  
jvandrun@uci.edu

<sup>1</sup> Logic and Philosophy of Science, UC Irvine, Irvine, CA 92687-5100, USA

purpose they use a language consisting of the words “block”, “pillar”, “slab”, “beam”. *A* calls them out;—*B* brings the stone which he has learnt to bring at such-and-such a call. Conceive this as a complete primitive language. (1958, §2)

In using the language, one agent calls out the words and the other acts on them. To know how to do so is to know the rules of a game. More generally, “[w]e can ... think of the whole process of using words in [the builder-assistant language] as one of those games by means of which children learn their native language.” The language game associated with the primitive builder-assistant language is “the whole, consisting of language and the actions into which it is woven” (1958, §7).

Inasmuch as it involves learning how to use the language, a language game on Wittgenstein’s conception is an evolutionary game. As the players repeatedly interact, they update their strategies on the basis of what has happened in past plays. They are playing the game well when their language use facilitates successful action.

Our aim here is not to reconstruct Wittgenstein’s philosophical views generally nor his account of how one might learn an established language more specifically. We are concerned, rather, with how a language game, where words are inextricably interwoven with action, might emerge from prelinguistic interactions between potential language users. This forging of a language game might also be characterized as an evolutionary game. It is a game that shows how language and action might come to form an integrated whole in the first place. In particular, we will consider how this might happen in the context of generalized signaling games that allow for self-assembly.

David Lewis (1969) introduced the idea of a signaling game to show how linguistic conventions might be established. Brian Skyrms (2006) subsequently showed how the classical games Lewis described might be reformulated as evolutionary games that illustrate how even low-rationality learners might evolve simple signaling languages. Barrett and Skyrms (2017) have more recently shown how both simple and more generalized signaling games might self-assemble by means of ritualized interactions. Better understanding the self-assembly of games is a key motivation for the present models.<sup>3</sup>

A self-assembling game is characterized by what the players do when they interact and how their structure of interactions evolves. The sender in a signaling game observes nature, then sends a signal. The receiver waits for the signal, observes the signal, then acts. Then the agents each update their dispositions based on what happened and do it all again. This sequence of interactions forms a network that is manifest across time. It is a structure where players take turns. The theory of self-assembling games is concerned with the self-assembly of such diachronic structures.

The generalized signaling games we consider here illustrate how a meaningful language and the diachronic structure of the players’ interactions might coevolve to facilitate successful action.<sup>4</sup> The first game shows how a player might learn to initiate

<sup>3</sup> See also Barrett et al. (2018), and Barrett (2020). See Barrett et al. (2020) and Steinert-Threlkeld (2020) for two accounts of how nontrivial linguistic compositionality might evolve in the context of signaling games under reinforcement learning.

<sup>4</sup> Skyrms (2010, pp. 154–155) describes a dialogue game with a fixed diachronic structure. Here we are concerned with how a diachronic structure like that might coevolve with the evolution of meaningful discourse.

meaningful discourse by asking a question rather than immediately acting. The second game shows how agents may evolve to ask new questions with coordinated meanings that coevolve on repeated plays. We will also consider variants of each game that illustrate features of self-assembly. Together, these games show how language and the structure of discourse itself may come to be interwoven with action and how the diachronic structure of an evolutionary game may coevolve with the play of the game.

## 2 The emergence of discourse

We will start with something akin to Wittgenstein's builder-assistant game. It shows how agents might come to be involved in meaningful discourse using an evolved language instead of simply acting in the first place. It also shows how they might learn to end discourse and to act instead of talking. While this game is very simple, it allows for self-assembly. It is this that determines the structure of discourse between the two agents.

The *question game* begins with nature randomly determining whether the builder needs a slab or a block with unbiased probabilities. The assistant knows that the builder needs a slab or a block but does not know which. He may guess at what the builder needs and hand her a slab  $S$  or a block  $B$  at random, or he may produce a signal  $Q$  instead of immediately acting.<sup>5</sup>

If the assistant decides to guess and hands the builder what she needs given the current state of nature, then the round is a success. In this case, the assistant is rewarded in a way that may influence his future actions, something we will discuss in more detail in a moment. If the assistant hands the builder the wrong thing given the current state of nature, then the round is a failure. In this case, neither the assistant nor the builder are rewarded, and they may even be punished, something else we will discuss in a moment. So if the assistant hands the builder a block or slab, players succeed or fail and the current round of the game ends.

If the assistant decides instead to signal  $Q$ , then the builder observes the signal and replies with one of two signals  $A_0$  or  $A_1$ .<sup>6</sup> In this case, the assistant has a second chance to guess what the builder needs by handing her a slab  $S$  or a block  $B$ , or he may again signal  $Q$ . If the assistant chooses to hand something and hands the builder what she needs given the current state of nature, then the round is a success and both players are rewarded. If the assistant hands the builder the wrong thing, then the round is a failure. In this case, neither player is rewarded and both may be punished. If the

<sup>5</sup> The signal is just a state of nature produced by the assistant to which the builder has epistemic access and on which she might consequently condition her actions. The various signals we will consider are assumed to be distinguishable from each other so that the listener can act in a way that depends on the particular signal sent.

<sup>6</sup> We are concerned here with how the assistant might learn to initiate discourse. Importantly, a similar story might be told for how the builder might learn to continue discourse by replying to  $Q$ . In each case, the process is one where the structure of the players' interactions self-assembles through the ritualization of successful actions. How this might work for various aspects of the structure of a language game is a recurring theme in the present paper. See Barrett and Skyrms (2017) and Barrett (2020) for discussions of how self-assembly works more generally.

assistant chooses to signal  $Q$  again, then we will also count the round as a failure on the present version of the game.

Note that the question game may be played without any discourse at all if the assistant always just chooses to guess what the builder needs. In that case, one should expect the builder and her assistant to be successful about half the time as a result of blind luck. For the builder and her assistant to do better than chance, they must coevolve an integrated set of dispositions. Specifically, the assistant must learn to initiate discourse with  $Q$ , the builder must learn to use  $A_0$  and  $A_1$  to represent the current state of nature, and the assistant must learn to respond by handing the builder the building material that has come to be represented by the builder's signal. If the builder and her assistant evolve to be uniformly successful,  $Q$  must come to initiate discourse much as the question "What do you need?" might, and  $A_0$  and  $A_1$  must come to specify each of the two building materials: slab and block.

In order to characterize an evolutionary game, we need to say how the players learn from their experience. To this end, we will consider two learning dynamics: *simple reinforcement* and *reinforcement with punishment*.<sup>7</sup> One might represent how the players' dispositions evolve under each of these dynamics by considering a process of adding and drawing balls from urns. The urns and the types of ball each contains represent the degrees of freedom of the players' dispositions and the proportion of each type of ball in each urn represent their first-order dispositions to act on states of nature and signals. Some of their actions will serve to structure their discourse. The question game on simple reinforcement learning proceeds as indicated in Fig. 1 read from top to bottom. The events on each round of play are as follows.

QUESTION GAME (SIMPLE REINFORCEMENT):

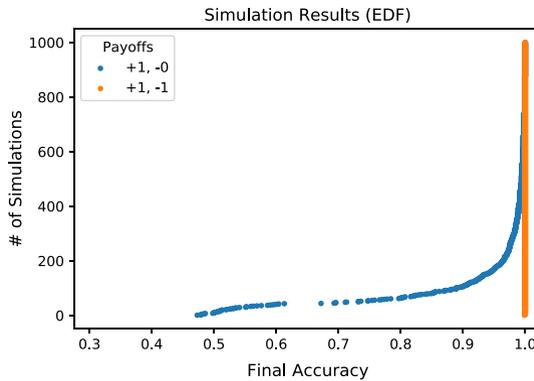
ASSISTANT MOVE I: Nature randomly determines whether the builder needs a slab or a block with unbiased probabilities. The assistant draws a ball from his *start urn*. This urn begins with one ball each of types  $Q$ ,  $S$ , and  $B$ . If he draws  $S$  or  $B$  he just hands the builder a slab or a block. If the builder gets what she needs given the state of nature, then the assistant returns the ball he drew to the urn from which he drew it and adds a duplicate ball of the same type. Else, the assistant just returns the ball he drew to the urn from which it was drawn. But if the assistant draws  $Q$ , he signals  $Q$ .

BUILDER MOVE I: If the assistant signals  $Q$ , the builder in turn draws a ball from an urn corresponding to the building material she needs. Specifically, if nature tells her that she needs a slab, then she draws from her *slab urn*; and if nature tells her that she needs a block, then she draws from her *block urn*. Each of these urns initially contains one ball each of types  $A_0$  and  $A_1$ . If the builder draws an  $A_0$  ball, she signals  $A_0$ ; and if she draws an  $A_1$  ball, she signals  $A_1$ .

ASSISTANT MOVE II: When the assistant hears the builder's reply, he draws from one of two reply urns,  $A_0$  and  $A_1$ , each initially containing one ball each of types  $Q$ ,  $S$ ,

<sup>7</sup> See Herrnstein (1970), Roth and Erev (1995), and Erev and Roth (1998) for a description of various forms of reinforcement learning and how they may be used to characterize human learning. See Barrett (2006) for a description and an early investigation of a number of learning dynamics and Barrett and Zollman (2009) for a discussion of the role of punishment and forgetting in learning. See Alexander et al. (2012), Barrett et al. (2017), Cochran and Barrett (2021), Cochran and Barrett (2022), Barrett and Gabriel (2021) for discussions of increasingly subtle learning strategies.





**Fig. 2** Final accuracies for the question game, displayed as an empirical CDF. Individual dots indicate the results of actual simulations, rank-ordered so that the corresponding value on the ordinate indicates the number of simulations out of 1000 which had less than or equal to the specified final accuracy. The blue distribution indicates the results for simple (+ 1, − 0) reinforcement, while the orange distribution the results for reinforcement with punishment (+ 1, − 1)

the assistant needs to know to hand the right thing to the builder given the current state of nature.

We have supposed that the builder simply gives up if the assistant asks a second question. A more subtle model would allow the conversation to continue at a cost for each question and reply. The cost might represent the time and effort the agents expend in the exchange.<sup>9</sup>

Consider the question game under simple reinforcement learning with a base reinforcement payoff of 2.0 for success and a fixed signal cost of 0.5 for each signal sent in the round. Here the assistant's expected return for randomly guessing is 1.0. This is less than the 1.5 expected return for asking one question if the builder and assistant have evolved a perfectly reliable signaling system. But it is more than the 0.75 expected return for asking one question when they have not yet evolved the ability to communicate.<sup>10</sup> So the question is whether the builder and assistant are able to evolve reliable enough signalling conventions fast enough to make discourse more attractive than guessing.

On simulation of the costly question game the builder gets what she needs with a final accuracy better than 0.75 on 0.534 of the runs and an accuracy of better than 0.9 on 0.398 of the runs with  $10^7$  plays per run. The assistant usually evolves to ask precisely one question then act on the builder's reply. This happens even in the case of suboptimal pooling as the expected return for asking one question using a language with a reliability of 0.75 is 1.125. On successful runs, the builder and assistant learn to signal reliably enough early on that the assistant is sufficiently likely to initiate discourse in future plays so that they have the chance to evolve a yet more reliable language which makes future discourse yet more likely. This feedback reinforces

<sup>9</sup> Costly signaling in nature is ubiquitous and well-studied. See Maynard Smith (1965), Sherman (1977), Zahavi and Zahavi (1997), and Searcy and Nowicki (2006) for discussions.

<sup>10</sup> For two questions, the assistant's expected return is the same as if he had just guessed up front; and for more questions, it is less. We will discuss what happens in such cases in a moment.

both the likelihood of future discourse and the accuracy of the evolved language. As a result, the agents coevolve a diachronic structure of interactions and a simple signaling language where the builder gets what she needs.

Unsurprisingly, the likelihood of discourse decreases as signaling costs increase. That said, reliable discourse is sometimes found to evolve on reinforcement learning in the present game with a signaling cost of 1.0. In this case, there is no expected advantage for the assistant to ask a question even after having evolved a perfectly reliable signaling system. Nevertheless, on simulation the builder gets what she needs with an accuracy better than 0.75 on 0.008 of the runs and with an accuracy better than 0.9 on 0.002 of the runs. If a reliable signaling system evolves quickly, the assistant may be regularly reinforced early on for initiating discourse in a round. As a result, the assistant may learn to communicate effectively with the builder and prefer doing so even when he might do slightly better in the medium run by guessing on his own.

To do better than chance in each version of the question game, the assistant must learn to initiate discourse and the builder and the assistant must together forge a language with sufficient expressive resources to tell the assistant what the builder needs. This requires them to coevolve dispositions that bind their signals and actions in such a way as to form an integrated whole that both structures their discourse and gives it content. Just as with Wittgenstein's language game, the game that self-assembles, when successful, allows for basic communication between the builder and her assistant.

The sequential structure of the language game they ultimately produce coevolves with the meanings of the players' signals. The options each player has on a round of play and the significance of those options depend on what the players did in earlier rounds, what happened when they did it, and how they updated their dispositions. The self-assembling game evolves a language game with a fixed diachronic structure together with the players' strategies for playing that game.

Important for the present argument, the agents cannot even begin to evolve a language with sufficient expressive resources for the task at hand without first getting involved in discourse. Self-assembly explains how the language part of the language game comes to be played at all. In brief, the agents talk because discourse *eventually* has salutary consequences. If they do not begin to play the language part of the game, they have no chance of evolving the stable linguistic practice that ultimately allows them to coordinate their practical actions. When they do start talking, it does not take long to get some measure of success. And this leads to further talk.

Also concerning the structure of discourse, the agents are only successful if the assistant learns when it is appropriate to stop asking questions and to act on the information shared by the builder. Knowing how to stop talking is as important as knowing how to start. But ending the conversation is only successful if the assistant gives the builder the right thing. It is in this way that the very structure of discourse co-emerges with meaningful language. Together they allow for successful action.

### 3 Meaning and the structure of discourse

In the question game, the assistant's one linguistic act evolves to serve as a prompt to get the builder to say something that might come to represent the material the builder needs. One can imagine a self-assembling *dialogue game* that allows for the evolution of more subtle language games.

Suppose that the builder needs one of four possible building materials on each round (red slab  $RS$ , red block  $RB$ , blue slab  $BS$ , or blue block  $BB$ ) and that the assistant has two potential linguistic acts ( $Q_0$  and  $Q_1$ ) that may come to represent the same or different questions. There are a number of ways to fill in the details to characterize a particular self-assembling game. We will discuss one of these in some detail, then briefly describe what happens in a natural variant.

Part of filling in the details involves saying what options and resources each player has and what might affect each player's actions at each step in the game. To begin, we will suppose that the builder has the same four responses available to answer each of her assistant's two potential questions ( $A_0$ ,  $A_1$ ,  $A_2$ , and  $A_3$ ) and that she conditions her actions on what she needs and the question that her assistant just asked and on nothing else. We will further suppose that the assistant conditions his actions on everything the builder has said so far in the round and on nothing else.<sup>11</sup>

The dialogue game on simple reinforcement proceeds as indicated in Fig. 3 read from top to bottom. The events on each round of play are as follows.

DIALOGUE GAME (SIMPLE REINFORCEMENT):

ASSISTANT'S MOVE I: Nature randomly determines what the builder needs with unbiased probabilities from among the four possible materials: red slab  $RS$ , red block  $RB$ , blue slab  $BS$ , or blue block  $BB$ . Her assistant then draws a ball from an urn that begins with one ball of each of the six types  $RS$ ,  $RB$ ,  $BS$ ,  $BB$ ,  $Q_0$ , and  $Q_1$ . If he draws  $RS$ ,  $RB$ ,  $BS$ , or  $BB$ , he simply hands the corresponding material to the builder. If it is what the builder needs, the round ends with success and the assistant returns his ball to the urn and adds a duplicate ball of the same type. If it is not what the builder needs, the play ends in failure and the assistant just returns the ball he drew to the urn. If the assistant draws  $Q_0$  or  $Q_1$ , he sends the corresponding signal.

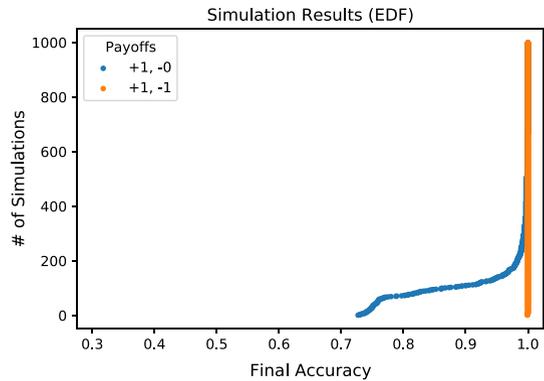
BUILDER'S MOVE I: The builder draws from an urn corresponding to her assistant's signal and the building material she needs. She has eight urns for this purpose labeled  $Q_0RS$ ,  $Q_1RS$ ,  $Q_0RB$ , ... each initially containing one ball each of four types  $A_0$ ,  $A_1$ ,  $A_2$ , and  $A_3$ . The builder sends the signal corresponding to the type of ball drawn.

ASSISTANT'S MOVE II: The assistant observes the builder's reply and draws a ball from one of four new urns  $A_0$ ,  $A_1$ ,  $A_2$ , and  $A_3$ , each corresponding to one of the possible replies. Each reply urn begins with one ball of each of the six types  $RS$ ,  $RB$ ,  $BS$ ,  $BB$ ,  $Q_0$ , and  $Q_1$ . If he draws  $RS$ ,  $RB$ ,  $BS$ , or  $BB$ , he hands the corresponding material

<sup>11</sup> One might consider games where the players condition on anything to which they have epistemic access. What they in fact find salient evolves by means of a learning dynamics that serves to ritualize whatever behavior leads to successful action. See Barrett and Skyrms (2017) and Barrett (2020) for details regarding how this works.



**Fig. 4** Final accuracies for the dialogue game



Again, an analogous description characterizes the dialogue game under  $(+1, -1)$  reinforcement with punishment. Success works just as described above, but failure leads to each ball drawn in the round being discarded unless it was the last ball of its type in the urn from which it was drawn.

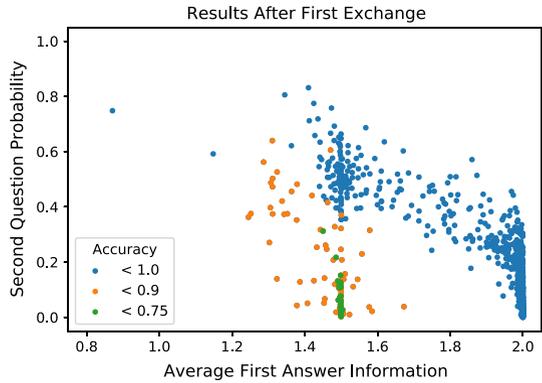
In order to be optimally successful, the builder and assistant must learn to communicate within an evolved structure of interactions using questions and answers that have coevolved coordinated meanings. The assistant must learn to initiate discourse, to ask the right questions given their evolving meanings, and eventually to stop asking questions. And the builder must evolve answers to the two questions the assistant may ask that together communicate what she needs.

Note that only the first question and answer are in principle needed in this initial version of the game. That said, asking both questions may still serve a purpose if the first question does not end up eliciting a response that fully specifies what the builder needs. If the agents do evolve to play beyond the first question on a round, the two questions might come to mean different things that together determine the assistant to provide the builder with the material she needs. Indeed, this often happens under simple reinforcement learning.

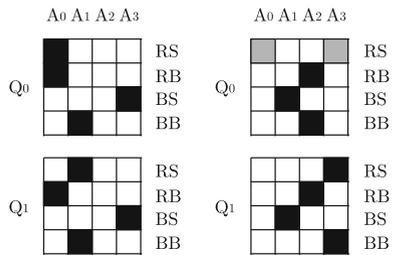
On simulation the players initially signal and act randomly, but on repeated plays, they again typically evolve nearly optimal dispositions. With  $10^7$  plays per run on simple reinforcement, they end up with actions that are reliable more than 0.9 of the time on 0.890 of the runs. On  $(+1, -1)$  reinforcement with punishment, all of the runs were observed to yield a nearly optimal final reliability. The blue distribution in Fig. 4 represents the number of simulation runs out of 1000 that did less than or equal to the specified final accuracy value on simple reinforcement. The orange does the same for  $(+1, -1)$  reinforcement with punishment. A significant proportion of the runs on simple reinforcement 0.072 exhibited a final success rate of less than 0.8. These suboptimal pooling equilibria can be seen in the blue distribution's inflection point at 0.75 in Fig. 4.

The agents are so successful in evolving an efficient language game on the  $(+1, -1)$  reinforcement with punishment game, that there is not much to say. Here the players evolve a highly accurate language that only requires one question. In contrast, language

**Fig. 5** The probability that the assistant asks a second question given the information (in bits) communicated by the first question. Once again, individual dots indicate actual simulations, with the blue dots corresponding to those simulations with final accuracies greater than 0.9, the orange dots to those with final accuracies between 0.9 and 0.75, and the green dots to those with final accuracies less than 0.75



**Fig. 6** Two examples of evolved languages (left and right) in which each question elicits insufficient information but together precisely specifies the material needed



games with different diachronic structures may evolve on simple reinforcement. The evolved diachronic structure of a game depends in part, then, on the adaptive dynamics.

On different runs of the simple reinforcement game, the players sometimes evolve linguistic practice that does not require a second question. The assistant needs two bits of information to know what material the builder needs. As Fig. 5 shows, the probability of the second question being asked at all in the language game resulting from a full run decreases as the evolved information content of the first answer increases.

Note, however, that sometimes both questions get asked even when the first answer is fully informative. Indeed, the second question sometimes evolves to mean precisely the same thing as the first, as indicated by the answers it elicits and the subsequent actions it produces. There are a number of ways that this may happen. If the meaning of the first question were slower to evolve on a run than that of the second, such redundancy may have played a role in early success on the run. Since asking a second question is cost free in the present game, the evolution of this sort of redundancy is unsurprising. One would expect it to occur less frequently if there were a cost to asking a second question. In this spirit, there is a very high cost to asking a third question in the present game, and it comes to be almost never asked on either of the two dynamics.

More interesting, the two questions sometimes evolve different meanings and only allow for optimal success by dint of the systematically interrelated meanings of the questions and the replies they elicit. Two examples of this can be seen in Fig. 6. Both of these are from runs that produced language games that allow for nearly optimal success.

In the first run (left), the builder uses  $A_0$  for both  $RS$  and  $RB$  in reply to  $Q_0$ , and she uses  $A_1$  for both  $RS$  and  $BB$  in reply to  $Q_1$ , but her two answers together serve to determine the required building material.<sup>12</sup> In the second run (right), the builder sometimes uses  $A_0$  and sometimes uses  $A_3$  for  $RS$  and uses  $A_2$  for both  $RB$  and  $BB$  in reply to  $Q_0$ , and she uses  $A_3$  for both  $RS$  and  $BB$  and does not use  $A_0$  at all in reply to  $Q_1$ , but again her two answers together serve to determine the required building material.<sup>13</sup>

The two questions here provide two chances for the agents to evolve a reliable system of communication on simple reinforcement. Sometimes one works and the other is not needed. And sometimes they both work and one sees a form of redundancy. This is akin to what provides success on a learning dynamics like reinforcement with invention.<sup>14</sup> But in the present game there is also another phenomenon at work. The agents sometime evolve to ask different, but coordinated, questions.

Consider a costly version of the dialogue game where the players may continue to talk indefinitely with a base reinforcement payoff of 3.0 for success and a cost of 1.0 for each question and reply. On simulation the assistant almost always evolves to ask precisely one question then to act on the builder's reply. The assistant's expected return by guessing here is 0.5, and his expected return after asking one question is 2.0 if the assistant and builder are able to evolve a perfect language. So given what we saw in the costly question game, it is not surprising that the assistant learns to initiate discourse. When he does so and when the builder and assistant are able to evolve a working language, they are also forging a diachronic structure of interactions. Future plays of the game will involve less guessing and more extended discourse. Again, the diachronic structure of the game coevolves with what happens when the agents interact on repeated plays.

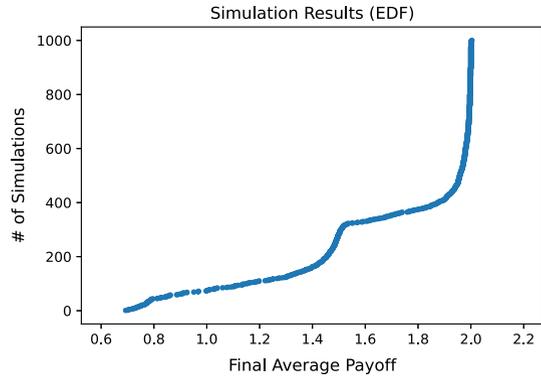
That the assistant does not evolve to ask more than one question also makes sense on reflection. The players usually 0.619 evolve a signaling language that gets the builder what she needs better than 0.9 of the time with  $10^7$  plays per run. But even when they get stuck in a suboptimal pooling equilibrium in the evolution of their language, the assistant only asks one question. At a suboptimal pooling accuracy of 0.75, he still has an expected return of 1.5, which is better than the expected return of 1.0 that he would get if he asked two questions even with a perfect language. So the builder and assistant are doing well enough playing a suboptimal pooling equilibria that the assistant does not have any reason to continue the conversation beyond one question. The distribution of payoffs in Fig. 7 illustrates this feature of the results. The upshot is that the assistant asking precisely one question is the optimal strategy if the builder and assistant are able to evolve a language with any real expressive power. And this is what the players are always found to do under the reinforcement dynamics.

<sup>12</sup> If she needs  $RS$ , she answers  $A_0$  to  $Q_0$  and  $A_1$  to  $Q_1$ . If she needs  $BB$ , she answers  $A_1$  to  $Q_0$  and  $A_1$  to  $Q_1$ . And for  $RB$  and  $BS$  she answers both questions  $A_0$  and  $A_3$  respectively.

<sup>13</sup> If she needs  $RS$ , she answers  $A_0$  or  $A_1$  to  $Q_0$  and  $A_3$  to  $Q_1$ . If she needs  $BB$ , she answers  $A_2$  to  $Q_0$  and  $A_3$  to  $Q_1$ . And for  $BS$  and  $RB$  she answers both questions  $A_1$  and  $A_2$  respectively.

<sup>14</sup> On this dynamics, new terms are invented during play. Since each new term has a chance of evolving a reliable meaning under simple reinforcement, each increases the chance that the signaling game will be successful. But they dynamics also produces a lot of redundancy in the form of evolved synonyms. See Alexander et al. (2012).

**Fig. 7** Final average payoffs for the costly dialogue game



Asking just one question is in principle enough for perfect success on the games we have discussed so far. But more than one question is always required for success in a slightly modified version of the dialogue game.

Consider a cost-free version of the dialogue game but where the builder only has the expressive resources for two possible answers  $A_0$  and  $A_1$  to each of the assistant's potential questions. On simulation the agents sometimes evolve optimal linguistic practice on both learning dynamics in this game as well. Here the only way that the builder can fully specify which of the four materials she needs is if her assistant asks both questions and if the answers to the two questions come to be systematically interrelated by means of evolved crosscutting properties. On successful runs the assistant learns to ask precisely two questions and the builder's reply to each question comes to pick out one of two coevolved cross-cutting kinds such that the two answers together serve to specify the precise building material needed.<sup>15</sup>

The various dialogue games illustrate how the meanings of the assistant's questions and the builder's replies might coevolve with the assistant's learning to ask questions then learning to stop asking when he has enough information to act successfully. The players self-assemble the structure of discourse and action that is needed for success given the evolving meanings of the terms. And they tune the meanings of their terms to the evolving structure of the game that they are playing. In this way, the resulting language game comes to exhibit a systematically interrelated whole, that ties together the evolved meanings of their words and the structure of their discourse and actions.

## 4 Discussion

Wittgenstein used the notion of a language game to illustrate how language is interwoven with action. We have shown here how a systematically interrelated whole, where

<sup>15</sup> This is similar to the natural-kind games described in Barrett (2007a, b) except, most saliently, here the game self-assembles to involve two questions as the meanings of the answers coevolve to represent conventional kinds. This modified dialogue game is significantly more difficult than these earlier natural-kind games or the dialogue games discussed above. On simple reinforcement in the present game, only 0.162 of the simulations have final accuracies greater than 0.75, with 0.078 having final accuracies greater than 0.9. On (+1, -1) reinforcement with punishment, 0.548 of the simulations reached final accuracies greater than 0.9.

the agents' words and the structure of their discourse and actions are thoroughly integrated to facilitate successful cooperative action, might be forged in the context of a simple learning dynamics. This explains how a simple language game like that described by Wittgenstein might come to be.

Self-assembly is essential to the agents' success in each of the games we have considered. The agents cannot even begin to evolve a language without first getting involved in discourse. And they cannot benefit from having evolved a language that allows for reliable communication without learning when to stop talking and use what they have learned.

The self-assembly of such diachronic structure is central to the theory of self-assembling games. It explains how the dispositions that allow for sequential interactions between agents might coevolve with other aspects of the game. It shows how agents may learn to take turns even as they coevolve what they will do on their turns.

Importantly, the present games illustrate only a part of the full self-assembly process. Just as the assistant learns to ask questions, the builder may learn to reply to the assistant's questions rather than to remain silent. To investigate how such a feature of discourse might evolve, one would give the builder the option of not replying at all or replying with one of a set of responses and see what happens under the adaptive dynamics. And so on for other structural features of the games we have considered. In short, there are yet more subtle models to investigate.

Such stories are thoroughly pragmatic. The agents learn to take part in discourse because it allows them to evolve a linguistic practice that eventually facilitates success in action. Achieving some measure of success leads to further talk. And that talk leads to further success. The feedback ritualizes their linguistic practice as a piece with their pragmatic practice more generally.

Ritualization of successful action under the adaptive dynamics of the self-assembling game is what structures the interactions of the agents and determines the significance of their actions at every step. It is what explains how they might come to use language at all.

The self-assembly of increasingly subtle language games allows for richer forms of meaningful discourse. In each evolved game, one's language and actions are inextricably interwoven.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

Alexander, J. M., Skyrms, B., & Zabell, S. L. (2012). Inventing new signals. *Dynamic Games and Applications*, 2, 129–145.

- Barrett, J. 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>
- Barrett, J. A. (2007a). Dynamic partitioning and the conventionality of kinds. *Philosophy of Science*, 74, 527–546.
- Barrett, J. A. (2007b). The evolution of coding in signaling games. *Theory and Decision*, 67(2), 223–237.
- Barrett, J. A. (2020). Self-assembling games and the evolution of salience. Forthcoming in *British Journal for the Philosophy of Science*.
- Barrett, J. A., & Gabriel, N. (2021) Reinforcement learning with punishment. manuscript.
- Barrett, J. A., Cochran, C. T., Huttegger, S., & Fujiwara, N. (2017). Hybrid learning in signaling games. *Journal of Experimental & Theoretical Artificial Intelligence*, 29(5), 1–9.
- Barrett, J. A., & Skyrms, B. (2017). Self-assembling games. *The British Journal for the Philosophy of Science*, 68(2), 329–353.
- Barrett, J. A., Skyrms, B., & Cochran, C. (2020). On the evolution of compositional language. *Philosophy of Science*, 87(5), 910–20.
- Barrett, J. A., Skyrms, B., & Mohseni, A. (2018). Self-assembling networks. *The British Journal for the Philosophy of Science*, 70(1), 301–325. <https://doi.org/10.1093/bjps/axx039>.
- Barrett, J. A., & Zollman, K. (2009). The role of forgetting in the evolution and learning of language. *Journal of Experimental and Theoretical Artificial Intelligence*, 21(4), 293–309.
- Cochran, C. T., & Barrett, J. A. (2021). How Signaling Conventions are Established. *Synthese*, 199(1–2), 4367–4391.
- Cochran, C. T., & Barrett, J. A. (2022). The Efficacy of Human Learning in Lewis-Skyrms Signaling Games, manuscript.
- Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88, 848–81.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243–266.
- Lewis, D. (1969). *Convention*. Harvard University Press.
- Maynard Smith, J. (1965). The evolution of alarm calls. *American Naturalist*, 99, 59–63.
- Roth, A. E., & Erev, I. (1995). Learning in extensive form games: Experimental data and simple dynamical models in the intermediate term. *Games and Economic Behavior*, 8, 164–212.
- Searcy, W. A., & Nowicki, S. (2006). *The evolution of animal communication: Reliability and deception in signaling systems*. Princeton University Press.
- Sherman, P. W. (1977). Nepotism and the evolution of alarm calls. *Science*, 197, 1246–1253.
- Skyrms, B. (2006). Signals. *Philosophy of Science*, 75(5), 489–500.
- Skyrms, B. (2010). *Signals: Evolution, learning, & information*. Oxford University Press.
- Steinert-Threlkeld, S. (2020). Toward the emergence of nontrivial compositionality. *Philosophy of Science*, 87(5), 897–909.
- Wittgenstein, L. (1958). *Philosophical investigations*. Basil Blackwell Ltd.
- Zahavi, A., & Zahavi, A. (1997). *The handicap principle: A missing piece of Darwin's puzzle*. Oxford University Press.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.