

LANGUAGE GAMES AND THE EMERGENCE OF DISCOURSE

JEFFREY A. BARRETT AND JACOB VANDRUNEN

ABSTRACT. Ludwig Wittgenstein (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.

1. INTRODUCTION

Ludwig Wittgenstein was concerned with how meaningful language was interwoven with action. 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).¹ 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 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

Date: August 5, 2021.

¹The Wittgenstein references are to the numbered sections of his 1958 *Philosophical Investigations*.

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.²

The generalized signaling games we consider here illustrate features of self-assembly. The self-assembly allows for the structure of the players' interactions itself to evolve on repeated plays. These games show how (1) a meaningful language and (2) the structure of discourse in that language might coevolve to facilitate successful action. The first game shows how players might learn to initiate meaningful discourse by asking a question rather than immediately acting. The second game and a variant we will briefly consider show how agents may evolve to ask new questions with coordinated meanings that coevolve on repeated plays.³ These games explain how language and the structure of discourse itself may come to be interwoven with action.

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 .⁴

²See also Barrett, Skyrms, and Mohseni (2018), and Barrett (2020). See Barrett, Skyrms, and Cochran (2020) and Steinert-Threlkeld (2020) for two accounts of how nontrivial linguistic compositionality might evolve in the context of signaling games under reinforcement learning.

³The second game might be thought of as a more general version of a game Skyrms described (2010, 154–5). On that game, the players are faced with different decision problems that require different information to solve. Nature presents one player with the decision problem. That player sends a signal to her friend who selects one of two partitions of nature to observe, then sends a reply. The first player then chooses an action to perform based on the return signal. This game illustrates how questions might evolve to make use of a fixed partition. In contrast, we show here how coordinated questions, partitions, and the structure of discourse itself might coevolve.

⁴At least initially, a 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.

If the assistant decides to guess and hands the builder what she needs given the current state of nature, then the interaction 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 interaction 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 is over.

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 .⁵ 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 interaction is a success and both players are rewarded. If the assistant hands the builder the wrong thing, then the interaction is a failure. In this case, neither player is rewarded and both may be punished. If the assistant chooses to signal Q again, then the interaction is a failure and both players may be punished.

Note that the question game may be played without any discourse at all if the assistant always chooses just 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 learning* and *reinforcement with punishment*.⁶ 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

⁵We 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.

⁶See Herrnstein (1970), Roth and Erev (1995), and Erev and Roth (1998) for a description of various forms of reinforcement learning and how they maybe used to characterize human learning in a number of salient contexts. See Barrett (2006) for 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, Skyrms, and Zabell (2012), Barrett, Cochran, Huttegger, and Fujiwara (2017), Cochran and Barrett (2020) and (2021), and Barrett and Gabriel (2021) for discussions of increasingly subtle learning strategies.

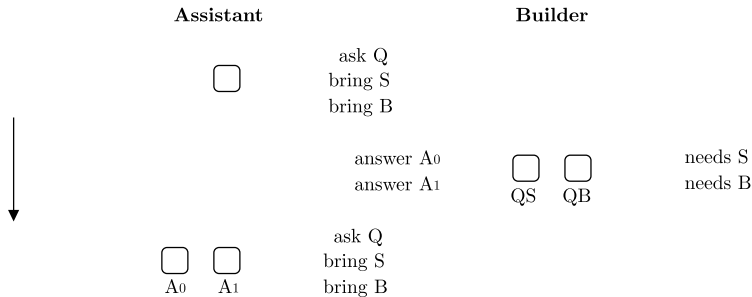


FIGURE 1. A question game. Play begins in the top left and zig-zags towards the bottom right. The far right column of text represents nature's play, while the middle column represents signals and/or actions taken by players. The boxes represent urns which the respective players draw from, conditioned on nature's play and/or the signals received.

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 Figure 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 , and B , then either signals Q again, or hands a slab S or block B to the builder. If the builder gets the building material she needs, then the round was successful and both players return the balls they drew to the urns from which they drew them and add a duplicate ball. Else, each agent just returns his or her balls to the urns from which they were drawn.

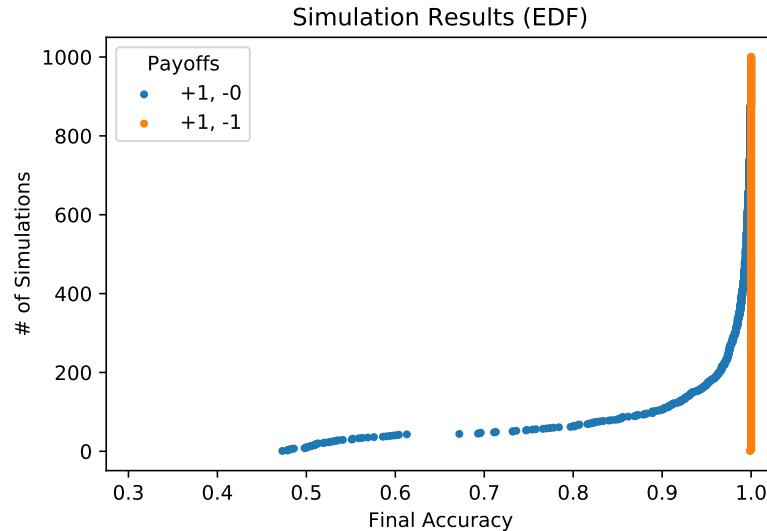


FIGURE 2. Final accuracies for the question game, displayed as an empirical CDF. Individual dots indicate the results of actual simulations, rank-ordered such that the corresponding value on the ordinate indicates the number of simulations out of 1000 which had less than or equal to the specified final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

A closely analogous description characterizes the question game under reinforcement with punishment. The difference is that here when a round leads to a successful action, each of the players returns the ball that he or she drew to the urn from which it was drawn and adds a copy of that ball. And when a round leads to failure, each of the players discards the ball that he or she drew unless it was the last ball of its type in the urn; in which case, he or she just returns the drawn ball to the urn from which it was drawn.

On simulation, the builder and her assistant begin by acting randomly, but on repeated plays, the assistant typically evolves to ask what the builder needs rather than just guess, and the builder's reply coevolves to represent what she needs. With 10^7 plays per run, on simple reinforcement, the players end up with dispositions that are reliable more than 0.9 of the time on 0.895 of the runs. For $(+1, -1)$ reinforcement with punishment, all of the runs were observed to yield a final reliability better than 0.9.

Figure 2 provides a more more detailed sense of these results. The blue distribution represents the number of simulation runs out of 1000 where the final accuracy was less than or equal to the specified value on simple reinforcement learning. The

dots indicate the actual final accuracy of each run. The orange distribution does the same for $(+1, -1)$ reinforcement with punishment.⁷

Inasmuch as Q , A_0 , and A_1 are initially meaningless, they are neither questions nor answers. But as the players interact with each other and the world and update their dispositions in accord with their learning dynamics, Q comes to serve as a question and A_0 and A_1 come to serve as answers, answers that communicate precisely what the assistant needs to know to hand the right thing to the builder given the current state of nature.

To be successful, 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 here allows for basic communication between the builder and her assistant.

Here the structure of the language game they ultimately produce coevolves with the meanings of the players’ signals.⁸ 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 structure together with the players’ strategies for playing that game.

Important for the present line of 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 start talking 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.⁹

Also concerning the structure of discourse, the agents are only successful in the present game 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.

⁷There are 1000 runs on each game. The final accuracy is determined by running each game 10^3 additional non-learning plays after the initial 10^7 and tracking how well the players do just in this last 10^3 plays. There is nothing special about this cutoff other than that it represents an accuracy higher than any observed suboptimal pooling equilibrium.

⁸A classical game or standard evolutionary game has a fixed interaction structure. A self-assembling game allows for both the structure and payoffs to evolve on repeated plays. Such structural evolution sometimes forges a classical game.

⁹In the present game, there are no barriers to such exploration since asking the first question and replying to it are both cost-free. A natural extension of this game would allow for costly signaling. What happens in that case would depend both on the level of the costs and on precisely how they were incurred. We will return to this issue briefly in the context of the dialogue game in the next section.

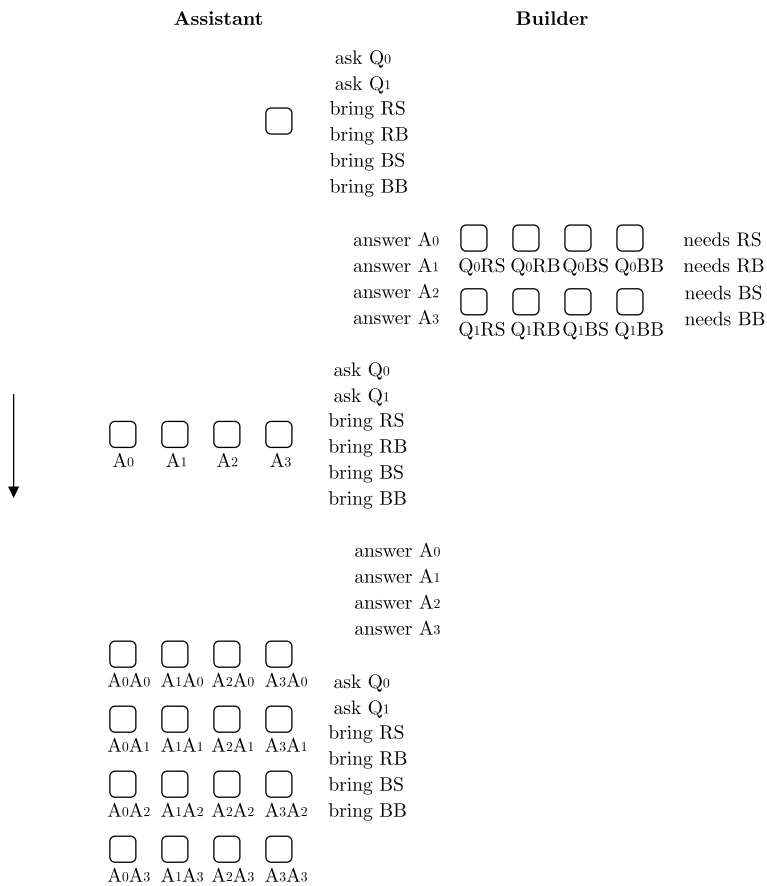


FIGURE 3. A dialogue game.

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 easily 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.¹⁰

The dialogue game on simple reinforcement proceeds as indicated in Figure 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 , \dots 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 to the builder. If the assistant gives the builder what she needs, the round

¹⁰One 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.

ends with success and both players return the balls they have drawn to the urns from which they were drawn and add a duplicate ball of the same type. If it is not what is needed, the round ends in failure and each just returns the ball(s) to the urns from which they were drawn. Else, if the assistant draws Q_0 or Q_1 , he sends the corresponding signal.

BUILDER’S MOVE II: This is exactly the same as builder move (i). The builder has only one set of urns, so the dispositions that the builder uses to reply to the first question are precisely the same as the dispositions she uses to reply to the second question.

ASSISTANT’S MOVE III: The assistant observes the builder’s second reply and draws from one of sixteen new urns depending on both the builder’s first and second replies labeled A_0A_0, A_0A_1, \dots . Each of these reply urns 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 gives the corresponding material to the builder, and the round ends in success or failure depending on whether the builder got what she needs. The players reinforce on all draws that led to success in the round and do not reinforce on failure. If the assistant draws Q_0 or Q_1 , he asks another question and the round ends in failure as the builder has lost patience. The players do not reinforce.

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 the present 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 act and signal 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 Figure 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 Figure 4.




FIGURE 4. Final accuracies for the dialogue game.

The agents are so successful in evolving an efficient language game on the $(+1, -1)$ reinforcement with punishment self-assembling game, that there is not much to say. In contrast, quite different language games may evolve on the simple reinforcement self-assembling game.

On different runs of the simple reinforcement game, the players sometimes evolve language practice that does not require a second question. The assistant needs two bits of information to know what material the builder needs. As Figure 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.

Perhaps 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 Figure 6. Both of these are from runs that produced language games that allow for nearly optimal success.

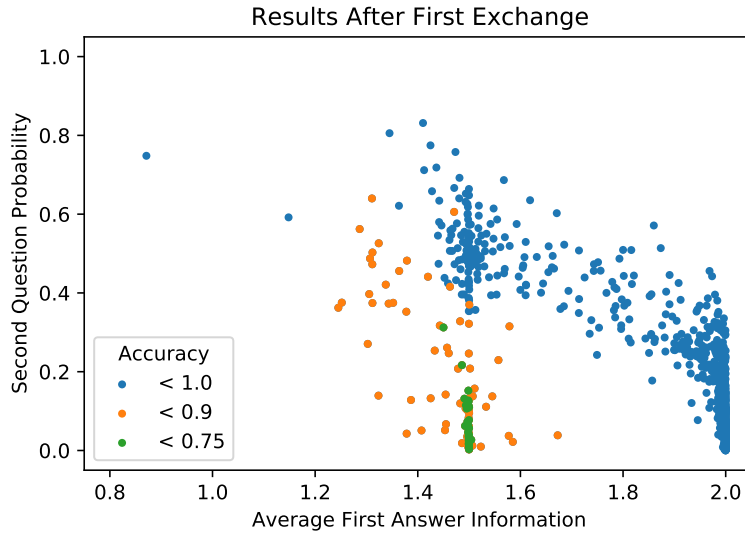


FIGURE 5. The probability that the assistant asks a second question given the information (in bits) communicated by the first question. Once again, individual dots indicate actual simulations, with the blue dots corresponding to those simulations with final accuracies greater than 0.9, the orange dots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

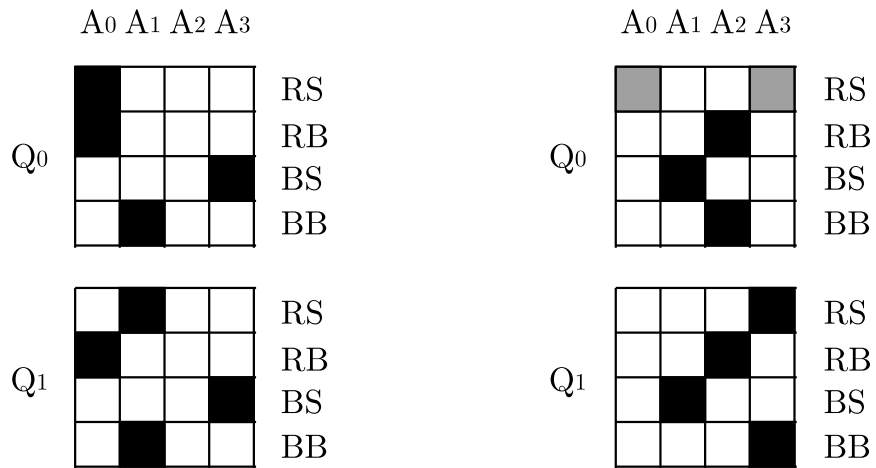


FIGURE 6. Two examples of evolved languages (left and right) in which each question elicits insufficient information but together precisely specifies the material needed.

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.¹¹ 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.¹²

The two questions here provide two chances for the agents to evolve a reliable system of communication on simple reinforcement. Sometimes one works and 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.¹³ But here there is also another phenomenon at work. The agents sometime evolve to ask different, but coordinated, questions.

While two questions are both essential to the agents' ultimate success on some runs of the present game, one is enough in principle. But asking two questions is always required for success in a slightly modified version of the game.

Consider a game like the one we just discussed 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. On successful runs, the assistant learns to ask precisely two questions and the builder's reply to each questions 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.¹⁴

In both the original and the modified dialogue games, the meanings of the assistant's questions and the builder's replies coevolve with the assistant's learning to ask questions then 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.

¹¹If 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.

¹²If 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.

¹³On 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, Skyrms, and Zabell (2012).

¹⁴This is similar to the natural-kind games described in Barrett (2007a) (2007b) 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 Barrett's natural-kind games or the dialogue game just discussed. On simple reinforcement, 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.

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 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.

Importantly, the games here illustrate only a part of the full self-assembly process. Just as the assistant learns to ask and to stop asking questions in the present games, 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 learning dynamics. And so on for other aspects of the agents' evolving structure of interactions.

What ultimately drives the evolution of discourse is the ritualization of successful action under the learning dynamics of the self-assembling game. It is this that structures the interactions of the agents and determines the significance of their actions at every step of their coevolved linguistic practice. It is this that explains how the agents might come to play a language game 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.

REFERENCES

- [1] J. M. Alexander, B. Skyrms, S. L. Zabell (2012). Inventing new signals. *Dynamic Games and Applications*, 2, 129–145.
- [2] Barrett, J. A. (2020). Self-assembling games and the evolution of salience. Forthcoming in *British Journal for the Philosophy of Science*.
- [3] Barrett, J. A. (2007a). Dynamic partitioning and the conventionality of kinds. *Philosophy of Science*, 74, 527–546.
- [4] Barrett, J. A. (2007b). The evolution of coding in signaling games. *Theory and Decision*, 67(2), 223–237.
- [5] Barrett, J. A. and Brian Skyrms (2017). Self-assembling games. *The British Journal for the Philosophy of Science*, 68(2), 329–353
- [6] Barrett, J. A., Brian Skyrms, and Aydin Mohseni (2018). Self-assembling networks. *The British Journal for the Philosophy of Science*, 70(1), 301–325. doi 10.1093/bjps/axx039
- [7] Barrett, J. A., Brian Skyrms, Calvin Cochran (2020). On the evolution of compositional language. *Philosophy of Science* 87(5): 910–20.
- [8] Barrett, J. A., C. T. Cochran, S. Huttegger, N. Fujiwara (2017) Hybrid learning in signaling games. *Journal of Experimental & Theoretical Artificial Intelligence* 29(5): 1–9.
- [9] Barrett, J. A. and N. Gabriel (2021) Reinforcement learning with punishment. draft.
- [10] Barrett, J. A and K. Zollman (2009). The role of forgetting in the evolution and learning of language. *Journal of Experimental and Theoretical Artificial Intelligence*, 21(4), 293–309.
- [11] Cochran, C. T. and J. A. Barrett (2021) “The Efficacy of Human Learning in Lewis-Skyrms Signaling Games,” draft.
- [12] Cochran, C. T. and J. A. Barrett (2020) “How Signaling Conventions are Established,” forthcoming in *Synthese*.
- [13] Erev, I. and A. E. Roth (1998) Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88: 848–81.
- [14] Herrnstein, R. J. (1970) On the law of effect. *Journal of the Experimental Analysis of Behavior* 13:243–266.
- [15] Houser, Nathan and Christian Kloesel (eds.) (1992). *The Essential Peirce: Selected Philosophical Writings Volume 1 (1867-1893)*. Bloomington and Indianapolis: Indiana University Press.
- [16] Lewis, David (1969) *Convention*. Cambridge, MA: Harvard University Press.
- [17] Roth, A. E. and I. Erev (1995) Learning in extensive form games: experimental data and simple dynamical models in the intermediate term. *Games and Economic Behavior* 8: 164–212.
- [18] Skyrms, Brian (2010). *Signals: Evolution, Learning, & Information*, New York: Oxford University Press.
- [19] Skyrms, Brian (2006). Signals. *Philosophy of Science*, 75(5), 489–500.
- [20] Steinert-Threlkeld, Shane (2020) Toward the emergence of nontrivial compositionality. *Philosophy of Science* 87(5): 897–909.
- [21] Wittgenstein, Ludwig (1958) *Philosophical Investigations*, GEMÁnscombe translator, Oxford: Basil Blackwell Ltd.