

ON THE COEVOLUTION OF BASIC ARITHMETIC LANGUAGE AND KNOWLEDGE

JEFFREY A. BARRETT

ABSTRACT. Skyrms-Lewis sender-receiver games with invention allow one to model how a simple mathematical language might be invented and become meaningful as its use coevolves with the basic arithmetic competence of primitive mathematical inquirers. Such models provide sufficient conditions for the invention and evolution of a very basic sort of arithmetic language and practice, and, in doing so, provide insight into the nature of a correspondingly basic sort of mathematical knowledge in an evolutionary context. Given traditional philosophical reflections concerning the nature and preconditions of mathematical knowledge, these conditions are strikingly modest.

First of all, it has to be noted that mathematical propositions, strictly so called, are always judgments a priori, not empirical; because they carry with them necessity, which cannot be derived from experience.

Immanuel Kant, *The Critique of Pure Reason*, B 14–15
Norman Kemp Smith translation

1. INTRODUCTION

There is a venerable and oft-honored tradition of considering mathematical knowledge to be somehow a priori, apodictic, universal, certain, and/or necessary. All human judgement, however, is a product of human nature, which, on our best understanding, is the result of contingent evolutionary processes. The problem then is to understand the nature of mathematical knowledge in a contingent evolutionary context.

Here we consider how it might be possible for a very basic sort of mathematical language, knowledge, and judgment to coevolve in the context of Skyrms-Lewis

sender-receiver games with invention and simple reinforcement learning.¹ Such models show *how it is possible* for simple agents to invent a basic mathematical language and coevolve dispositions that track mathematical truth. In doing so, they provide insight into the nature of a basic type of mathematical knowledge in an evolutionary context.²

It is important to have the line of argument clear from the start. We will consider how a descriptive number language might be invented and a basic sort of arithmetic competence might coevolve in a concrete dispositional model. More specifically, in such models, the terms of the descriptive number language are invented by the agents, then coevolve with the agents' developing arithmetic competence to be meaningful. We will then consider what the evolved language and mathematical practice concern in such models. We will find that the primitive arithmetic language and competence of the modeled agents most directly reflects (1) the nature of the world they inhabit and (2) the second-order dispositions they use to revise their first-order dispositions to signal and to act in that world. This provides an explanation then for how it is possible for relatively simple agents to coevolve a basic sort of mathematical language and knowledge.³

For convenience, we will refer to the terms that evolve in the context of the evolutionary games as *number terms*. The precise meaning of the language the agents evolve, however, is determined by the precise details of the coordinated action that their evolved language supports, which is, in turn, given by the nature of the world the agents inhabit and the agents' second-order dispositions. On most of the games discussed here, when successful, the agents coevolve terms that represent the cardinality properties of finite collections together with dispositions that represent a competence in cardinal addition. The suggestion then will be that one should expect a correspondingly richer mathematical language and more

¹Such games were first used by David Lewis[9] to account for the formation of convention. Brian Skyrms later described how such games might be formulated in an evolutionary context without any assumptions of ideal rationality or common knowledge[14]. The type of invention-learning dynamics we will use for the arithmetic games was proposed by Brian Skyrms [13]. Some of its formal properties are further considered by Skyrms, Alexander, and Zabell [1]. It is an extension of the sort of simple reinforcement learning proposed by Roth and Erev [12].

²Skyrms has shown how dispositions that track basic logical truths might evolve in simple Skyrms-Lewis games [15]. Here we consider rather more involved games that allow for the invention of mathematical language and the evolution of mathematical dispositions. See [10] for a recent survey of evolutionary accounts of other sorts of convention.

³In particular, the agents start with no language whatsoever and only random first-order dispositions to act on signals. As we will see, the subtlety of the evolved practices of such agents depends on the precise nature of their second-order dispositions and the world they inhabit.

sophisticated varieties of arithmetic competence to coevolve when agents are faced more subtle arithmetic tasks.

2. SENDER-RECEIVER GAMES AND THE EVOLUTION OF LANGUAGE

In the simplest variety of Skyrms-Lewis sender-receiver game there are two agents: a sender and a receiver.⁴ The sender observes the state of nature, then sends a signal. The receiver, who cannot observe nature directly, observes the signal then performs an action that either matches the state of nature and is successful or does not match the state of nature and is unsuccessful. If the act is successful, then the disposition that led to each agent's last action is reinforced; otherwise, it is not reinforced and may be weakened. Indeed, all it means for an action to be successful in the context of such a game is that it generates a result that, given the nature of the world that the agents inhabit and the agents' second-order dispositions to update their first-order dispositions to signal and to act, leads to the reinforcement of those first-order dispositions that produced the action; similarly, all it means for an action to be unsuccessful is that, given the nature of the world and the agents' second-order dispositions, it generates a result that does not lead to the reinforcement of the first-order dispositions that produced the action.⁵

The signals are meaningless when the agents begin to play the game, but as the sender's dispositions to signal (conditional on the state of nature) and the receiver's dispositions to act (conditional on the the sender's signal) evolve, the sender's signals become meaningful insofar as they serve as the basis for successful coordinated action. Whether an ideally successful signaling system evolves in such a game depends on the number of states of nature, the statistical distribution of states of nature, the agents' signaling resources, their initial first-order dispositions

⁴For discussions of standard Skyrms-Lewis sender-receiver games and other variants see [9], [14], [2], [6], [5], [4], and [13]). While we will restrict our attention here to evolution in the context of learning models, one should expect similar results in the context of population models as there is, for example, a close formal relationship between evolution under reinforcement learning and evolution under the replicator dynamics.

⁵While useful, the distinction between the agents' first-order dispositions (their dispositions to signal and to act) and their second-order dispositions (the learning dynamics that updates their first-order dispositions) is at best a rough one. It would begin to unravel a bit if, as is certainly the case in the evolution of real languages, the agents' second-order dispositions were themselves allowed to evolve in response to the evolution of their first-order dispositions. In particular, one should expect real inquirers to better learn how to learn as they evolve better descriptive language and the practices that this better language supports. The rough distinction between first- and second-order dispositions, then, is perhaps best understood as one between faster-evolving dispositions that primarily concern signaling and acting and slower-evolving dispositions that primarily concern updating the faster-evolving dispositions.

to signal and to act, and the second-order learning dynamics that update these initial dispositions.

A concrete example of a four-state sender-receiver game is represented in Figure 1 below. Here there are four uniformly distributed states of nature the sender might observe, four signals the sender might use, and four actions the receiver might perform. The states of nature might be, for example, (0) clear day, (1) clear night, (2) rain, and (3) snow. And the possible receiver actions might be (0) bring sunglasses, (1) bring flashlights, (2) bring umbrellas, and (3) bring snowshoes. If the receiver's action matches the state of nature, then the agents are successful and their actions are reinforced; otherwise they are unsuccessful and not reinforced. More specifically, we will suppose, in this case, that the agents learn by bounded reinforcement with punishment.

Here one might characterize the agents' first-order dispositions by the contents of various types of balls in urns and a description of how the contents of the urns determines the agents' actions and characterize the agents' second-order dispositions by how the contents of the urns are updated given the natural events that result from their actions.

Regarding the agents' first-order dispositions, we will suppose that the sender has one urn for each possible state of nature labeled 0, 1, 2, and 3. Each of her urns begins with one ball of each type of signal she might send labeled A , B , C , and D . The receiver has one urn for each type of signal. Each of his urns begins with one ball of each type of action he might take labeled 0, 1, 2, and 3. The sender observes the weather, draws a ball at random out of the corresponding urn, with each particular ball equally likely, then sends the signal on that ball. The receiver observes the signal, draws a ball at random out of the urn corresponding to the signal, with each particular ball equally likely, then performs whatever weather-related action is indicated on the ball drawn.

Regarding their second-order dispositions, we will suppose here that if the receiver's action matched the weather, then the sender and receiver put the ball they drew back in the urn they drew it from and add to that urn another ball of the same type if there are fewer than N_{\max} balls of that type already in the urn. This will typically strengthen the associated dispositions. On the other hand, if the receiver's action did not match the weather, then each agent puts back the ball he drew if and only if it was the last ball of its type, and discards it otherwise. This will typically weaken the associated dispositions. Note that, given these second-order dispositions, there will always be at least one ball of each type and never more than N_{\max} balls of any particular type in each urn.




FIGURE 1. A simple four-state sender-receiver game.

Such agents are easily simulated. The sender and receiver begin by randomly signaling and randomly acting. Any initial success they have is consequently the result of blind luck. As their first-order dispositions evolve under the influence of their second-order dispositions, however, the agents nearly always evolve a set of nearly optimally successful, and systematically interrelated, dispositions to signal and to act. Before the simulation is run, it is impossible to say which term will evolve to mean what, but each term typically does in fact evolve to represent a different state of nature and hence together they allow the agents to match the state of nature with its corresponding action. That they have evolved a meaningful language is manifest in their evolved success.⁶

Since the sender observes the state of nature and since the receiver acts for both agents, the agents must evolve dispositions to faithfully represent the states of nature in their signals and systematically related dispositions to act appropriately on the basis of the signals in order to be successful.

⁶More specifically, for $N_{\max} = 1000$ and 1000 runs each with 10^6 plays, the sender and receiver nearly always (0.993) evolve a set of nearly optimal (0.994) dispositions. Similar results are obtained for a wide range of parameter values for this general sort of bounded reinforcement learning with punishment. With unbounded reinforcement and no punishment, the agents in this game evolve an optimal signaling language only about 0.78 of the time and end up in a suboptimal mixed equilibrium the other 0.22 of the time.[6] What leads to the very high degree of success in the case above is that (i) there is a maximum number of each type of ball in each urn (so bad habits, if they arise, do not get too strongly ingrained), (ii) no type of signal or action ever goes to extinction in any urn (so there is always a possible escape from suboptimal dispositions), and (iii) both reinforcement and punishment, or weakening, of first-order dispositions are possible (since without reinforcement, the agents would be unable to learn; and without punishment they would be unable to forget suboptimal dispositions and hence would lack the means of escape). With this learning dynamics the fact that no action goes to extinction allows for the possibility of the agents' first-order dispositions randomly wandering away from success, but, if they do, they will quickly return and spend most of their time almost ideally successful in their actions.

3. THE INVENTION OF NUMBER TERMS AND THE COEVOLUTION OF BASIC ARITHMETIC COMPETENCE

Consider the more subtle sender-receiver game illustrated in Figure 2 below. This game involves a king and three agents: two senders, the king's right- and left-hand advisors; and one receiver, the royal jeweler who makes gold rings for the court.⁷

The game is played as follows. The king holds out fingers on his left and right hands. Being an arbitrary monarch, his dispositions are such that each possible configuration of fingers is equally likely. The king's left-hand advisor looks at the king's left hand and his right-hand advisor looks at his right hand. They each then send a signal to the royal jeweler, who makes a collection of rings for the king. If there is precisely one ring made for each finger that the king extended, then the king is pleased and rewards his agents, which, in turn, affects their future dispositions. But before saying precisely how, it is important to be clear concerning how the agents' actions are determined.⁸

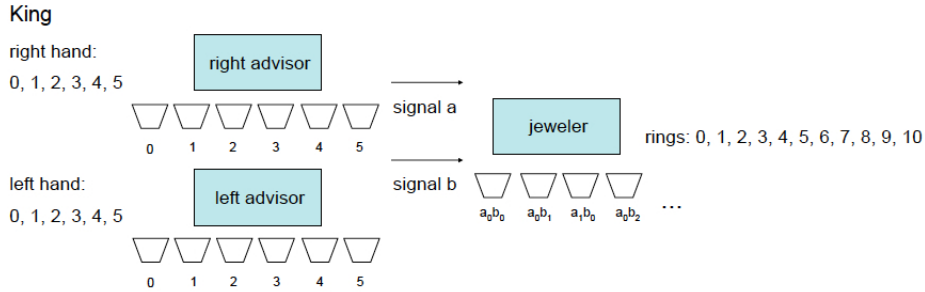


FIGURE 2. A sender-receiver game that evolves basic arithmetic.

The king's advisors determine what signal to send to the jeweler using balls and urns, but differently than above. Here each advisor has six urns, each initially containing just a single black ball, and each affixed with a picture of the king's hand, one displaying zero fingers, another displaying two fingers, etc., to five fingers. While there is always precisely one black ball in each of the advisor's urns, other balls, labeled with terms corresponding to potential signals, are added to the urns

⁷Note that the king is not an agent in the game; rather, he is simply part of the world that the agents inhabit. Of course, how nature responds to the acts of the agents determines in part what stable first-order dispositions, if any, will evolve in the game.

⁸Since the agents in this game are both evolving a language and evolving a simple theory of arithmetic, as represented in their evolved dispositions, this game might alternatively be characterized as a sender-predictor game. See [4] for a description of such games.

on repeated plays of the game. More specifically, when the king holds out his fingers, each advisor goes to the urn that matches the fingers on the hand that he sees, then draws a ball at random from that urn. If the ball is black, the advisor invents a new term at random and sends that term to the jeweler; alternatively, if the ball is labeled with a term, the advisor sends that term to the royal jeweler.⁹ On each play of the game, then, the jeweler receives a pair of terms, one from each of the king's advisors.

The jeweler decides what collection of gold rings to make for the king as follows. The jeweler has an urn labeled by each pair of terms that he has received from the advisor so far, and he constructs further new urns as needed as the advisors invent and send new terms. Each of the jeweler's urns begins with one ball affixed with a picture of no rings, one ball affixed with a picture of one ring, etc., up to one ball affixed with a picture of ten rings.¹⁰ When the jeweler gets the pair of signals from the king's advisors, he goes to the urn corresponding to the two signals and draws a ball at random. He makes one gold ring for each ring on the picture on the drawn ball, then delivers the rings to the king.¹¹

⁹This simple rule for inventing terms was proposed by Skyrms [13]. There are, at least in the short run, more effective ways to introduce new terms. One might, for example, tie the likelihood of the introduction a new term to the failure of old terms in promoting successful action. Such a rule might also more strongly discourage the unnecessary invention and evolution of synonyms. That said, if agents can evolve successful behavior using Skyrms' very simple invention rule and reinforcement without punishment, this only serves to illustrate the robust nature of the evolution toward successful practice in such games.

¹⁰More precisely, when he constructs a new urn, he populates it with one ball of each act type allowed in the game as it is currently being played. Here this always means one ball for each of the eleven possible actions: making 0 to 10 rings. The condition *as it is currently being played* will matter later when we consider games where the king changes the structure of the game during play.

¹¹While the agents do not need prior pure concepts, forms of sensible intuition, intentions, or such to play this game, they do need second-order dispositions that allow them to learn. In the case of urn learning by simple reinforcement, among other things, this means that the agents must be able to match states to urns and signals to urns consistently. Having such second-order dispositions is simply one of the preconditions for the game. To be sure, one might expect that such second-order dispositions are themselves often the product of prior evolutionary processes. In particular, one might imagine that the ability to match objects in a one-to-one onto way is a good candidate for an evolved ability. And, indeed, the four-state game in the last section shows how it is possible for agents to evolve dispositions that jointly represent a one-to-one onto map between states, terms, and acts where there was no such map to begin. The general point here, however, is more important. What language and practice evolves on a particular game *if any* depends on the precise second-order dispositions of the agents playing the game and the nature of the world they inhabit.

Nature and the agents' second-order dispositions affect future plays of the game as follows. The king puts the rings delivered by the royal jeweler on his extended fingers. If there is precisely one ring for each extended finger, then the advisors and the jeweler return the ball they drew to its urn and also add one ball corresponding to the action they took to that urn. If an advisor drew a black ball and the jeweler's action was successful, then he also adds one ball of the new signal type he invented to each of his urns, and the jeweler adds urns corresponding to the new types of signals he may receive. If the jeweler's action was not successful, if the agents did not match gold rings to royal fingers, they just return the balls they drew to the urns.¹²

While the overall task before the king's advisors and jeweler is significantly more difficult than the task faced by the agents in the four-state game of the last section, the way in which the agents learn here is in one important sense less sophisticated. Once a new term has been invented, the king's agents learn by simple reinforcement without punishment, arguably the least sophisticated learning dynamics possible.

Such agents are easily simulated. Since the king's advisors do not initially have any terms to send, they begin by inventing new terms, and at a relatively high rate since there is not much in their urns to keep them from drawing the black ball often. The advisors initially send their newly-minted terms at random, and the jeweler initially acts randomly when he receives a new pair of terms. As the agents update their first-order dispositions by reinforcement, however, they typically evolve a set of nearly optimal, systematically interrelated dispositions where the advisors look at the king's hands, send terms, and the jeweler makes precisely the number of rings as the number of extended fingers on both of the king's hands.¹³ The terms sent by the advisors, then, evolve to represent numbers (cardinality properties) and the jeweler's dispositions evolve such that he *adds* those numbers (cardinal addition) in order to determine how many rings to make.¹⁴ In this sense, then, the king's advisors and the royal jeweler typically coevolve a basic arithmetic competence with

¹²Note that a new type of term is only kept if the first-play of the term was successful and that it then becomes available for any purpose.

¹³In 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. Because there is no punishment, it is invention that does the work here in avoiding suboptimal equilibria. When a modest degree of punishment or forgetting is added to invention, the simulated agents typically evolve yet more quickly and surely toward an optimal equilibrium between their first- and second-order dispositions.

¹⁴More specifically, in this particular game, the number terms are invented, then coevolve with the arithmetic competence of the agents to represent the cardinalities of concrete collections.

number terms that successfully represent the states of nature and dispositions to act on these terms that reflects basic arithmetic truth.¹⁵

While it is unclear how much arithmetic might be coevolved, invented, and/or discovered by such primitive means, there is reason to believe that the longer one plays such a game, the greater the expressive power of their language and the more basic arithmetic knowledge exhibited by the agents.

In order to see this more concretely, consider a more sophisticated version of the last addition game. Suppose that a king, with an unbounded number of fingers, begins with simpler arithmetic problems where he displays at most k fingers on each hand, then proceeds to more complex problems where he displays at most $k + 1$ on each hand when his advisors have reached a sufficiently high success rate (say 0.95) on the simpler problems.¹⁶ And keep the same learning dynamics with simple reinforcement and invention.

On simulations of this graduated addition game, the agents coevolve meaningful number terms and the coordinated dispositions to add numbers (cardinality properties) for a small initial k , they extend this basic arithmetic competence to the new $k + 1$ problems, and quickly and reliably invent, evolve, and/or discover a significant portion of second-grade arithmetic, including the philosophically salient fact that $7 + 5 = 12$. Significantly for the question at hand, they appear to do ever better the longer they play, providing modest empirical evidence that their growth in arithmetic competence is unbounded.¹⁷ And, given the concrete model,

¹⁵Since it is always possible to draw a black ball, the number of terms used by the advisors increases over time, but at an ever slower rate. Since there is no cost for keeping terms in use and no mechanism for forgetting in this model, numerous synonyms evolve over time. Further, since there is nothing in the model to coordinate the meanings of the terms used by the two advisors, each advisor evolves his own number language. The jeweler, then, must not only learn to add the numbers signaled by the advisors but also to coordinate between the synonyms in each advisor's language and between the advisors' different languages. And the jeweler typically evolves to do precisely this, and the three agents together coevolve to exhibit the basic arithmetic competence. When successful, the agents have coevolved the ability to represent the cardinalities of two finite collections of concrete objects and to represent the cardinality of the collection that would be formed by combining those two collections.

¹⁶Note that this game has no equilibria since nature keeps changing as the agents evolve; or, equivalently, one might put it that the game that the agents are playing itself keeps changing, preventing the agents from reaching equilibrium. The agents here are nevertheless typically able to track the moving target of successful action. This tracking provides a more general notion of success than the notion of approaching a stationary equilibrium between first- and second-order dispositions.

¹⁷On 1000 runs of 1.0×10^6 plays each, the mean difficulty k of the sums the agents are able to do with accuracy 0.95 is 8.03; on 2.0×10^6 plays, they are at mean difficulty of 10.27; and with 4.0×10^6 plays, they are at mean difficulty of 13.43. When they have played as long as the agents

one can see just how they did it and what precisely their coevolved dispositions are tracking.

4. THE NATURE OF MATHEMATICAL KNOWLEDGE

Such models show *how it is possible* for relatively simple agents to coevolve basic number terms and dispositions that track simple arithmetic truths.¹⁸

They also show how it is possible for linguistic and arithmetic competence to coevolve from dispositions that initially represent nothing, to dispositions that represent vague meanings and approximate arithmetic truths, to dispositions that together represent sharp meanings and precise arithmetic truths. And how it is possible for the coevolved representation of basic arithmetic meanings and truths to be so closely intertwined that one cannot say which of the agents' dispositions represent meanings and which represent truths. A term that coevolves to represent 5 only gradually comes to represent a number (here a cardinality) at all, and then, only in the systematically intertwined dispositions of the agents. Similarly, a term that coevolves to mean 7 when sent by one advisor, means what it does because, among other things, it leads to the appropriate action of the jeweler making 12 rings when the other advisor sends a term that coevolves to mean 5. The representation

in the last game, 8.0×10^6 plays, they are typically successfully expressing and adding numbers up to $17 + 17$. And after 1.6×10^7 plays, the agents have typically evolved the language to express and the ability to compute sums reliably up to $22 + 22$ and after 3.2×10^7 up to $28 + 28$. In short, the arithmetic competence of the agents appears to grow significantly faster than linearly against the log of the number of plays. This provides modest reason to believe that the agents' arithmetic competence may be unbounded even in this relatively simple game. Of course, one would much prefer having a theorem. While Alexander, Skyrms, and Zabell have formal results that characterize the behavior of simple reinforcement games with invention[1], there is, so far, no proof that the agents in a game like the one described here would almost always evolve to exhibit arithmetic competence better than any given finite bound.

¹⁸Note that such models illustrate how it is possible for a very basic sort of mathematical competence to evolve given relatively weak preconditions, not how such competence evolved for human agents. Providing a how-in-fact explanation for even the simplest arithmetic competences exhibited by human agents would require one to know details about our evolutionary history, including features of our linguistic environment, that we will likely never be able to determine. If so, we are limited to consider only those philosophical purposes that may be served with much weaker how-possible explanations. Of course, providing a full how-possible explanation for the sophisticated sort of arithmetic competence exhibited by human agents would also be extraordinarily difficult. Since capturing more sophisticated mathematical practice requires more sophisticated evolutionary models, the thought is that the project best proceeds in clear, concrete steps. This approach also allows more sophisticated models to use abilities evolved in the simpler models.

of the arithmetic truth $7 + 5 = 12$ is similarly the gradual result of the coevolution of dozens of subtly interrelated dispositions.

In the context of such models, one can say precisely what the particular coevolved mathematical language and knowledge concerns. When the agents in fact manage to coevolve an arithmetic competence, their evolved competence represents *their second-order dispositions* to update their first-order dispositions to signal and to act *given what in fact happens* when they signal and act. The evolved arithmetic truth $7 + 5 = 12$, for example, most directly reflects the *de facto* payoffs, in the form of reinforcements, given the way the agents are built and the nature of the world they inhabit. Had the king rewarded other configurations of rings or had the king's rewards not influenced the agents as they did or had the world produced different events on the same agent actions, then the dispositions of the agents might have evolved to track something very different than the standard addition. More specifically, had the agents been equipped with different second-order dispositions, they might have evolved number terms, but number terms associated with equality, modular addition, multiplication, or absolute difference instead of ordinary elementary addition.¹⁹ And agents inhabiting a sufficiently different world or equipped with sufficiently different payoffs might not evolve terms that we recognize as representing numbers or dispositions that we recognize as representing standard mathematical truths.

The king's agents invent number terms and evolve simple arithmetic competence, when they do, then, because this is what is required for their first-order dispositions to evolve to satisfy the demands of their second-order dispositions given the nature of the world they inhabit. When their first-order dispositions so evolve, the agents have managed, on the endogenous standard of success represented by their second-order dispositions, to evolve successful first-order dispositions given the nature of the world they inhabit. To be sure, the evolution of a descriptive mathematical language and the dispositions corresponding to addition here involve the evolution

¹⁹Instead of learning how to match n rings and m rings to $n + m$ fingers and discovering standard addition, equipped with different payoffs, the agents might have learned how to match n rings to each of m fingers and discovered multiplication. The details of such evolutionary stories vary depending on the second-order dispositions one considers. Modular addition and absolute difference evolve, on simulation, very much as standard addition above. Multiplication evolves similarly but significantly slower since the number of possible outcomes of the operation, on at least one natural understanding, grow quickly with the number of possible arguments. On a slightly different type of game, the advisors might observe the king's fingers and the receiver might act on the signals and only be successful if the number of fingers on each hand were the same. Here the agents coevolve number terms and the relation of equality rather than an arithmetic operation. The type of number terms evolved might be thought of as correspondingly less subtle.

of conventions, but the agents' dispositions are not merely conventional as it is a simple matter of fact whether their evolved dispositions are successful given their endogenous mathematical standards.

There is then a primitive notion of normativity at work in such evolutionary models, and it is in the service of objective knowledge. Successful arithmetic inquiry here is evolution toward a reflective equilibrium between the inquirers' first-order dispositions to signal and to act and their second-order dispositions to reinforce their dispositions to signal and to act given the nature of the world they inhabit. The agents are evolutionarily successful on this endogenous standard precisely insofar as their first-order dispositions to signal and to act produce events that require their second-order dispositions to change their first-order dispositions ever less over time. Similarly, epistemic error in arithmetic inquiry occurs when the agents' first-order dispositions to signal and to act produce events that undermine stable action by requiring their second-order dispositions to continue to significantly revise their first-order dispositions. It is in this way that a primitive sort of distributed epistemic norm is exhibited in the agents' second-order dispositions.

While this is a relatively weak account of mathematical norms it is, at least, a clear one. Moreover, it is strong enough to require that the arithmetic inquirers in fact evolve terms that reliably communicate how many fingers the king holds out and that they in fact evolve the ability to correctly add the numbers so represented in order to be successful by the norms represented in their simple second-order dispositions.

5. CONCERNING EVOLUTIONARY ACCOUNTS OF MORE SOPHISTICATED VARIETIES OF MATHEMATICAL KNOWLEDGE

The evolutionary models described here explain how it is possible for relatively simple agents to invent meaningful number terms and to evolve dispositions that represent basic arithmetic truths. In doing so, such models provide a weak, but perfectly clear, account of what it might mean for a primitive sort of mathematical inquiry to be successful.

The number terms that evolve in such models represent properties of states of nature and allow the senders to signal the receiver to act successfully on states to which he has no access otherwise. The coevolution of this signaling language and the associated arithmetic dispositions involves the establishment of conventions, but these conventions are in the service of the endogenous norms represented in the second-order dispositions of the agents and there is a matter of fact whether they are successful or not. In particular, these norms can only be satisfied when

what actually happens when the agents act no longer leads to a revision of their first-order dispositions, which depends on objective matters of fact concerning the nature of the world they inhabit. The agents, then, are only successful when they get the world right, or at least that part of the world that matters to satisfy their second-order dispositions.

There are, of course, richer and more subtle forms of mathematical knowledge than those illustrated by these models. A richer variety of arithmetic knowledge might, for example, require that one have a *general* rule for adding.²⁰ Or that one *be able to prove* the truths of arithmetic from appropriately basic axioms.²¹ Or that one be able to account for the evolution of concrete judgments of mathematical necessity. Or even that one *somehow be able to directly picture or intuit* the intentional content of number terms and the truths of arithmetic as terms referring to objects in a complete number system.²² While the models described here show how it is possible for a basic sort of arithmetic practice to evolve, one would like to be able to do the same for more subtle varieties of mathematical practice, insofar as one can make clear sense of these.²³

²⁰While the agents modeled here learn a rule for addition and follow it, a more sophisticated sort of mathematical knowledge might require that agents learn a rule that would allow them to add numbers they may never have considered before [8]. In order to provide an evolutionary account of this sort of arithmetic competence, one would need to show how it is possible to coevolve first-order dispositions that represent a more general algorithm for addition than the direct concrete dispositions that evolve for the agents discussed here. See [3] for a discussion of the coevolution of language and more subtle sorts of rule-following behaviors that complement those discussed here. One would like to have a story, for example, for how terms evolved in one context might evolve to be used more generally. Two such evolutionary mechanisms are discussed in [3]. The simpler of the two is substitution. On substitution, when the task changes, say from adding the number of fingers on two hands to adding the number of marbles in two piles, rather than evolving a new coordinated descriptive language and dispositions for adding, it can be evolutionarily more efficient to evolve to use the old evolved dispositions in a new way. What one needs for the appropriation of the old evolved dispositions to a new task is for the new stimuli to come to trigger the old dispositions in a way that leads to successful action. A pile of five marbles and a pile of two marbles might then evolve to be treated the same way for the purposes of addition as five fingers on one hand and two on another.

²¹Providing a full account of how we give and evaluate proof would also require one to account for proofs that rely, for example, on geometric construction, which would require a very different sort of evolutionary model than the those discussed here.

²²The simple models discussed here clearly do not capture a rich notion of intentionality; rather, they show that a rich notion of intentionality is not required to account for the evolution of at least one very basic sort of arithmetic practice.

²³More sophisticated types of evolved mathematical language and practice can be expected to involve more subtle evolutionary games. The thought, again, is that the investigation of such games best proceeds in clear, concrete steps, where one may use the first-order dispositions evolved

One would expect such varieties of mathematical practice, understanding, and knowledge to be in the service of the second-order dispositions of the agents and interactions between agents and the world. The task is to provide an account of how it might be possible to evolve increasingly sophisticated practice on a set of relatively weak and compelling sufficient conditions. Insofar as this is possible, one will be able to say precisely what one means by rule-following, mathematical proof, mathematical necessity, or intentional content in the particular evolutionary context characterized by the model.

in simpler games to characterize the second-order dispositions of the agents in increasingly more subtle games.

REFERENCES

- [1] Alexander, J. M., B. Skyrms, S. Zbell (2011) “Inventing New Signals” *Dynamic Games and Applications* (2011): 1-17.
- [2] Argiento, Raffaele, Robin Pemantle, Brian Skyrms and Stas Volkov “Learning to Signal: Analysis of a Micro-Level Reinforcement Model,” *Stochastic Processes and their Applications* 119(2):373–390.
- [3] Barrett, Jeffrey A. (2012) “The Evolution of Simple Rule Following,” forthcoming in Brian Skyrms and Simon Huttegger edited volume.
- [4] Barrett, Jeffrey A. (2012) “On the Coevolution of Theory and Language and the Nature of Successful Inquiry,” forthcoming in Alistair Isaac edited volume.
- [5] Barrett, Jeffrey A. (2009) “Faithful Description and the Incommensurability of Evolved Languages,” *Philosophical Studies* 147(1): 123–137.
- [6] Barrett, Jeffrey A. (2007) “Dynamic Partitioning and the Conventionality of Kinds,” *Philosophy of Science* 74:527–546.
- [7] Bond, Alan B., Alan C. Kamil, Russell P. Balda (2003) “Social Complexity and Transitive Inference in Corvids,” *Animal Behaviour* 65: 479–487.
- [8] Kripke, Saul (1982) *Wittgenstein on Rules and Private Language: an Elementary Exposition*. Cambridge, Mass.: Harvard University Press.
- [9] Lewis, David (1969) *Convention*. Cambridge, MA: Harvard University Press.
- [10] Huttegger, Simon (2011) “How Much Rationality do We Need to Explain Conventions?” forthcoming.
- [11] Huttegger, Simon (2007) “Evolution and the Explanation of Meaning,” *Philosophy of Science* 74: 1–27.
- [12] Roth, Al and Ido Erev (1995) Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Intermediate Term *Games and Economic Behavior* 8: 164–212.
- [13] Skyrms, Brian (2010) *Signals Evolution, Learning, & Information*, New York: Oxford University Press.
- [14] Skyrms, Brian (2006) “Signals,” *Philosophy of Science* 75(5): 489–500.
- [15] Skyrms, Brian (2000) “Evolution of Inference,” in Tim Kohler and George Gumerman (eds.) *Dynamics of Human and Primate Societies*, New York: Oxford University Press, 77–88.