

On the Evolution of Compositional Language

Jeffrey A. Barrett, Calvin Cochran, and Brian Skyrms*

We present here a hierarchical model for the evolution of compositional language. The model has the structure of a two-sender/one-receiver Lewis signaling game augmented with executive agents who may learn to influence the behavior of the basic senders and receiver. The model shows how functional agents might coevolve representational roles even as they evolve a reliable compositional language in the context of costly signaling. When successful, the evolved language captures both the compositional structure of properties in the world and the compositional structure of successful actions involving those properties.

1. Introduction. Humans and some animals use languages that allow for the functional composition of basic terms to form more complex expressions. The meanings of the more complex expressions are influenced, sometimes determined, by the meanings of their parts. Animals for which there is compelling evidence for such compositional languages include putty nosed monkeys (Arnold and Zuberbühler 2006, 2008), Campbell's Monkeys (Ouattaraa, Lemassona, and Zuberbühler 2009), suricates (Manser, Seyfarth, and Cheney 2002), prairie dogs (Frederiksen and Slobodchikoff 2007; Slobodchikoff, Paseka, and Verdolin 2009), and some species of birds (Suzuki, Wheatcroft, and Griesser 2016).

The sort of composition exhibited by a language may be subtle. Suricates compose two acoustical aspects of their call to indicate predator class together with the signaler's perception of urgency. The compositional parts of the call are immediately evident to the suricates who hear it and can be detected by experimenters by means of an acoustical analysis of the call (Manser et al. 2002).

*To contact the authors, please write to: Jeffrey A. Barrett, 51 Schubert Ct., Irvine, CA 92612; e-mail: jabarret@uci.edu.

Correction: This article was reposted on December 24, 2020, to replace figure 4 and change "interested" to "uninterested" in footnote 1.

Philosophy of Science, 87 (December 2020) pp. 910–920. 0031-8248/2020/8705-0011\$10.00
Copyright 2020 by the Philosophy of Science Association. All rights reserved.

Broadly speaking, a language exhibits functional composition when the meanings of expressions in the language are functions of the meanings of their parts. But there are many ways that this might happen.¹ Here we are concerned with how simple compositional languages might evolve in the context of a *generalized signaling game*. The simplest sort of signaling game was introduced by Lewis (1969). It involves two players: a sender and a receiver. On each play of the game, nature chooses one of the states with unbiased probabilities and reveals it to the sender. The sender then sends a signal to the receiver, who cannot see the state of nature directly. The receiver chooses an act conditional on the signal. The agents are successful if and only if the act matches the current state of nature.

Treating this classical game as a repeated evolutionary game, the agents may then get a payoff based on the success of their action that affects their subsequent game-playing dispositions. Precisely how this works is specified by the learning dynamics one considers. Here we will consider simple reinforcement and reinforcement with costs. These are among the simplest learning dynamics possible and have been used to model both human and animal learning (Herrnstein 1970; Roth and Erev 1995; Erev and Roth 1998).

Lewis signaling games have been studied extensively under different learning and population dynamics, and there are a number of analytic and simulation results (Skyrms 2006; Barrett 2007; Hofbauer and Huttegger 2008; Argiento et al. 2009; Barrett and Zollman 2009; Skyrms 2010; Hu, Skyrms, and Tarrès 2011; Huttegger et al. 2014). Many variations of the basic signaling game have also been studied. These include games in which nature is biased, the agents have too few or too many signals, and there are multiple senders or receivers (Barrett 2006, 2013; Skyrms 2009).

The model we consider here is something we call a *hierarchical signaling game*. It has two basic senders and one basic receiver and an executive sender and an executive receiver. The basic agents play sender and receiver roles as in a basic signaling game, while the executive sender and executive receiver track and control aspects of the behavior of the basic agents. For the agents to be successful they must together evolve an efficient and reliable compositional language.

The basic senders have no preassigned representational roles in the game. Rather, the representational role each sender assumes is something that must evolve over time. As a result, the executive agents have the task of learning both what roles the basic agents have assumed and how their dispositions might be used for successful action even as the basic agents are evolving reliable dispositions. The dispositions of all of the agents then coevolve as the

1. We are uninterested in giving necessary or sufficient conditions for a language exhibiting compositional structure. We take different notions of compositionality to be suitable to different explanatory aims. The varieties of functional composition we consider here are closely allied with logical conjunction. Such composition is ubiquitous in natural languages.

compositional language evolves. On the current model, it is signaling costs that drive the agents toward efficient signaling and hence toward a compositional language.²

Costly signaling is ubiquitous in nature. In bacteria, and other microorganisms, each signal sent may involve producing a molecule that diffuses into the vicinity. Here each signal has a metabolic cost. For higher organisms, if giving an alarm call exposes an individual to increased immediate danger, then pausing to give two signals might increase danger. Again, this gives each signal a cost. Inefficient signaling may also carry costs in time or computational or attentional resources at both the sending and receiving end. As a result, costly signaling in nature is a well-studied phenomenon with an extensive literature (see, e.g., Maynard Smith 1965; Sherman 1977; Zehavi and Zahavi 1997; Searcy and Nowicki 2006). The costs in the current model covary with signal length. This might be taken to represent such things as attentive or computational costs in biological systems.

As with simpler signaling games, generalized signaling games might self-assemble from the ritualization of individual actions of the agents (Barrett and Skyrms 2017; Barrett, Skyrms, and Mohseni 2019). The behavior of the agents in the complex game is forged in the context of an evolutionary process. In the current model, this is a learning dynamics. As is often the case, the agents who play roles in the model here might be understood as the functional components of a single individual.

We briefly consider a basic signaling game and a simple generalized signaling game to set the stage. We then turn to the hierarchical composition game.

2. A Basic Signaling Game. The simplest Lewis signaling game involves a sender who observes one of two states and sends one of two signals and a receiver who performs one of two actions. On each play of the game, nature chooses one of the states with unbiased probabilities. The sender observes the state, then sends her signal. The receiver observes the signal, then performs his action. Each state of nature corresponds to exactly one act, and the players are both successful if and only if the act chosen by the receiver corresponds to the current state of nature. The evolutionary question concerns the conditions under which the agents might evolve a signaling system in which the sender associates exactly one signal with each state of nature, the receiver associates each of these signals with the corresponding act, and the players are always successful.

Simple reinforcement learning for this game can be pictured in terms of balls and urns. The sender has two urns, one for each state of nature (0 or 1), each beginning with one a -ball and one b -ball. The receiver has two urns, one

2. See Barrett (2007), Franke (2014), and Steinert-Threlkeld (2016) for other approaches to compositional signaling.




Figure 1.

for each signal type (a or b), each beginning with one 0-ball and one 1-ball. The sender observes nature, then draws a ball at random from her corresponding urn. This determines her signal. The receiver observes the signal, then draws a ball from his corresponding urn. This determines the act. If the act matches the state, then it is successful and each agent returns the ball to the urn from which it was drawn and adds a duplicate of that ball. If unsuccessful, then each agent simply returns the ball drawn to the urn from which it was drawn. In this way successful dispositions are made more likely conditional on the states that led to those actions (see fig. 1).

The sender begins by randomly signaling and the receiver by randomly acting, but if nature is unbiased, one can prove that this game will almost certainly converge to a signaling system in which each state of nature produces a signal that leads to an action that matches the state (Argiento et al. 2009). Of course, the evolved language here involves no composition at all. The evolution of compositional language requires a more subtle setup.

3. A Generalized Signaling Game. Consider a signaling game with two senders, one receiver, four states of nature, and four corresponding actions. Here nature chooses one of the four states of nature with unbiased probabilities. The two senders each observe the full state of nature. Each then randomly draws a 0 or 1 signal ball from her corresponding state urn. The receiver observes both signals and who sent them and then draws an act ball from his corresponding signal urn. If the act matches the current state, then it is successful and each agent returns the ball drawn to the urn from which it was drawn and adds a duplicate of that ball; otherwise, each agent just returns the ball drawn to the urn from which it was drawn.³

The senders again begin by sending random signals, and the receiver acts randomly. But here a simple compositional language typically evolves under simple reinforcement. On simulation, with 10^6 plays per run the agents are always observed to do better than chance, and they exhibit a cumulative success rate of better than 0.80 about 73% of the time (Barrett 2007). This level of success requires the senders to adopt systematically interrelated representational

3. This model is discussed in detail in Barrett (2007).




Figure 2.

roles, and it requires their evolved language to exhibit a corresponding sort of compositionality (see fig. 2).

When successful, the two senders coevolve to cross-partition nature in such a way that their two signals together represent the four states of nature. That is, each sender evolves to observe a particular conventional property in nature, and the properties each observes are systematically related in such a way that specifying whether each obtains constitutes a definite description that fully specifies the current state. The evolved language here is compositional. Specifically, each signal represents whether the corresponding property obtains, and the composition of the signals represents *simple conjunction*.

While the evolved language is compositional, conjunction is a very simple form of functional composition. And since the terms associated with each sender are never used alone, the language does not look very compositional. Rather, each pair of signals looks like a single signal in a simple game like the one described in the last section.⁴

To get a more subtle sort of composition, we need a more subtle game. Just as with the two games we have considered so far, the following hierarchical composition game might self-assemble from the ritualization of individual actions of the agents and modular composition of simpler games (Barrett and Skyrms 2017; Barrett et al. 2019).

4. A Hierarchical Composition Game. The hierarchical signaling game has two basic senders, one executive sender, one basic receiver, and one executive receiver. The basic senders and receiver play roles very similar to those of the agents in a standard Lewis signaling game, while the executive sender and receiver coevolve dispositions to track and to control the behavior of the basic agents. To be successful here, the agents must evolve a compositional language that allows for meaningful and efficient signaling.

4. See Franke (2014) for a critical discussion of this model.

In this game the state of nature features properties and a context. For concreteness, we consider two properties: color, which is either black or white, and animal, which is either dog or cat. Hence, the state of nature on a particular play of the game will feature black dog, black cat, white dog, or white cat.

The context indicates the type of information the receiver will need in order to perform the sort of action that will be successful on the current play. We consider three contexts: color, animal, or both. Figure 3 shows what the agents see, what urns they have to draw from given what they see, and what signals they might send or actions they might perform given what they draw.⁵

On each play of the game, nature chooses the state and the context randomly and with uniform probabilities. Both basic senders see the two natural state properties that obtain. The executive sender sees the context. Each basic sender draws from her property urn matching the current state to determine her signal, and the executive sender draws from her context urn matching the current context to determine which of the two basic senders (or both) will send her signal. Initially, all of these actions are random, but over time the basic senders may evolve roles and begin to communicate information about the natural state, and the executive sender may coevolve the ability to send the type of signal the current context demands.⁶

The basic receiver has four urns, one for each combination of sender *A* and *B* signals he might see. Each urn begins with one ball for each of the possible color-animal pairs: black dog, white dog, black cat, or white cat. The executive receiver has three urns, one for each combination of senders (*A*, *B*, or both) who may send signals. Each of these urns starts with one ball for each way the basic receiver might interpret the ball that he draws (color, animal, and both). The basic receiver sees the senders' signal(s), and then he draws from the urn corresponding to the sender *A* and sender *B* signal combination he sees if he sees a signal from both senders. If he only sees a signal from one sender, then he randomizes with uniform probability between the two relevant urns given the signal and who sent it. The executive receiver sees who signaled and then determines whether the basic receiver will interpret his draw as a color action, an animal action, or a composite action. The basic receiver then performs an action that corresponds to the ball he drew

5. See Barrett, Skyrms, and Cochran (2018) for two closely related compositional models and for further details regarding the current model.

6. More specifically, each sender is equipped with an urn for every possible state of nature, and each initially contains a 0-ball and a 1-ball. The executive sender is equipped with an urn for each of the three possible contexts: color, animal, and both. Each urn begins with one ball of each type sender *A*, sender *B*, and both. Upon witnessing the context, the executive sender randomly draws a ball from the corresponding urn. The drawn ball determines who will send a signal. Only the sender who is instructed by the executive sender to send her signal does so.

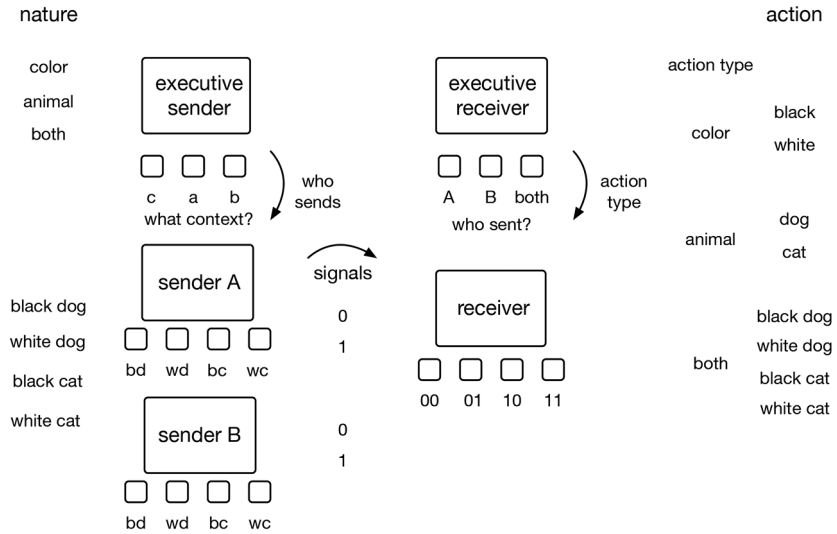


Figure 3.

given the interpretation specified by the executive receiver. For concreteness and to see how linguistic composition might correspond to composition in action, one might take the receiver’s action to involve producing an image on an initially blank canvas: black, white, dog, cat, black dog, white dog, black cat, or white cat.

We will suppose that the agents learn by simple reinforcement with a fixed cost per signal. This form of reinforcement learning involves three parameters: the simple-context payoff, the complex-context payoff, and the signal cost. When the basic receiver performs the act that matches the current state and context, then (1) if the context for that play is simple (color or animal), the reward is the simple-context payoff, and (2) if the context is both, the agents receive the complex-context payoff. And every signal sent has a cost. This represents the sender’s computational and representational costs and the receiver’s attentional and computational costs. For simplicity, we suppose that the costs per signal are the same for each. We suppose further that the reinforcements by which the sender and receiver learn are a simple function of the basic payoff and the signaling costs. Specifically, the fixed cost of each signal sent is simply subtracted from the payoff, and the resulting number of balls are added to (or subtracted from) the urns used in the play.

If the act does not match the state, then the basic payoff is zero, but since they still have signaling costs to pay, the agents will lose balls from the urns that led to the failed action. If the number of balls of a type in any urn goes below one, we set it to one. Agents are thus most strongly reinforced when they use the least number of signals to produce the action that matches the

current state and context, and they are punished if the signal costs they incur are higher than the basic payoff for their action given the context.

We measure success by tracking the proportion of the time the agents (1) chose a correct act given the state and context and (2) send the minimum number of signals needed to convey that act. The thought is that, regardless of how successful their actions may be, the agents have only evolved to use a compositional language if their behavior exhibits this sort of efficiency.

Consider a simple-context payoff of 1.5, a complex-context payoff of 2, and a signal cost of 0.5. On simulation with the strict standard of success just described, after 1,000 runs of 10^8 plays each, the mean overall cumulative success rate on simulation is 0.994. And of these runs, 97.5% exhibit near perfectly efficient signaling. Indeed, the agents usually exhibit near perfectly efficient signaling within 10^6 plays.

Success here is less sensitive to the relationship between the simple-context payoff and the complex-context payoff than it is to the relationship between the basic payoffs and signal cost. Setting the simple-context payoff and the complex-context payoff both to 2 and keeping a signal cost of 0.5 produces a slightly lower overall success rate of 0.973 with 93.6% of runs exhibiting near perfectly efficient signaling. This is still very good. But there is a significant drop in success if one lowers the signal cost relative to the two basic payoffs. When the simple-context payoff is 1.5, the complex-context payoff is 2, and the signal cost is 0.3, the agents evolve nearly perfectly efficient signaling 61% of the time. And when the signal cost is 0.1 they evolve near perfectly efficient signaling only 26% of the time.

Signal cost is responsible for efficient signaling and hence the form of functional composition exhibited. In order to be successful, the agents must both choose a correct act given the state and context and send the minimum number of signals needed to convey that act. Without a signal cost there is no evolutionary pressure on the number of signals the senders send on a play. And without that, there is nothing pushing in the direction of signaling efficiency and, hence, meanings for both composite expressions and the individual terms.

The languages that evolve here are compositional in a richer sense than in the model we discussed in the last section. Their evolved terms are typically used both together and alone. The meanings of the terms, when they evolve to have independent meanings, typically determine the meanings of the whole. And there is a clear sense in which the structure of the evolved compositional language mirrors both the structure of composition in nature and the sort of composition inherent in the receiver's successful actions.⁷

7. This last feature is a condition for rich linguistic composition that Josh Armstrong suggested in conversation.




Figure 4.

Further, composition on this type of hierarchical model is sometimes significantly more subtle than simple conjunction. As in figure 4, the ensuing sub-optimal language may have terms that fail to have meanings of their own yet play functional roles in modifying the meanings of other terms.⁸ While sender *B*'s terms refer to black and white, sender *A*'s terms do not individuate between different colors or different animals. Rather, the referents of *A*'s terms crosscut both animal and color to individuate between white dog or black cat and white cat or black dog. Given the salient contexts presented by nature, *A*'s terms are consequently entirely useless by themselves. This failure in efficiently marshaling the representational resources of the agents has the consequence that if the context requires just animal, the agents are helpless and cannot do better than chance. But sender *A*'s terms are useful when composed with sender *B*'s terms. When the context requires both animal and color, *B*'s term communicates a color and, together with *A*'s term, also selects an animal. It is not that *A*'s terms are meaningless. Rather, given the specified contextual demands, they are only useful in compositional structures.⁹

8. This particular language evolved in a run of the hierarchical model under reinforcement learning in which the payoffs were determined by the strict success condition above.

9. In natural language we routinely employ expressions that play important semantic roles in combination with other expressions but are relatively useless on their own. English language adverbs, adjectives, and pronouns often behave this way. While the term *only* is rarely useful on its own, expressions like *only child* or *only decaf* have precise meanings that may allow one to usefully characterize states of nature and, hence, facilitate successful action.

On the current model the compositional structure of the language often captures the prior compositional structure of the world—the structure of the properties we used to tell the narrative story of what the agents are learning to do. The compositional structure of the language also often evolves to represent the compositional structure of the operations involved in producing successful actions.

5. Morals. When the agents are successful in evolving a nearly perfectly efficient language on the current model, the basic senders evolve representational roles and coevolve signaling conventions appropriate to each role, the executive sender learns which roles each sender has adopted and how to use her signals to represent the current context, the executive receiver learns how to interpret the type of expression sent, and the basic receiver learns how to associate the signal and expression type with a successful action. Along the way, the meanings of the individual terms and the composite expressions coevolve to satisfy the expressive demands required for efficient communication.

That the signals have a cost plays a critical role here. When the agents evolve a system for nearly perfectly efficient signaling, both the individual terms and the composite expressions are meaningful, and the meanings of the composite expressions are a function of the meanings of the individual terms. As the language evolves, the semantic function of composition coevolves with the meanings of the individual terms. It is the expressive demands under the constraint of costly signaling that drives the evolution of the compositional language.

A hierarchical composition game like that described here might self-assemble by reinforcement. If it does, a language may evolve that has a relatively rich compositional structure.

REFERENCES

- Argiento, R., R. Pemantle, B. Skyrms, and S. Volkov. 2009. "Learning to Signal: Analysis of a Micro-Level Reinforcement Model." *Stochastic Processes and Their Applications* 119 (2): 373–90.
- Arnold, Kate, and Klaus Zuberbühler. 2006. "The Alarm-Calling System of Adult Male Putty-Nosed Monkeys, *Cercopithecus Nictitans Martini*." *Animal Behaviour* 72:643–53.
- . 2008. "Meaningful Call Combinations in a Non-human Primate." *Current Biology* 18 (5): R202–R203.
- Barrett, J. A. 2006. "Numerical Simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar." Paper no. 54, Institute for Mathematical Behavioral Sciences. <http://repositories.cdlib.org/imbs/54>.
- . 2007. "Dynamic Partitioning and the Conventionality of Kinds." *Philosophy of Science* 74:527–46.
- . 2013. "The Evolution of Simple Rule-Following." *Biological Theory* 8 (2): 142–50.
- Barrett, J. A., and B. Skyrms. 2017. "Self-Assembling Games." *British Journal for the Philosophy of Science* 68 (2): 329–53.

- Barrett, J. A., B. Skyrms, and C. Cochran. 2018. "Hierarchical Models for the Evolution of Compositional Language." Technical Report MBS 18-03, Institute for Mathematical Behavioral Sciences, University of California, Irvine.
- Barrett, J. A., B. Skyrms, and A. Mohseni. 2019. "Self-Assembling Networks." *British Journal for the Philosophy of Science* 70 (1): 301–25.
- 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.
- 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.
- Franke, M. 2014. "Creative Compositionality from Reinforcement." In *The Evolution of Language: Proceedings of the 10th International Conference*, ed. E. A. Cartmill, S. Roberts, H. Lyn, and H. Cornish, 82–89. Singapore: World Scientific.
- Frederiksen, J. K., and C. N. Slobodchikoff. 2007. "Referential Specificity in the Alarm Calls of the Black-Tailed Prairie Dog." *Ethology, Ecology and Evolution* 19: 87–99.
- Herrnstein, R. J. 1970. "On the Law of Effect." *Journal of the Experimental Analysis of Behavior* 13:243–66.
- Hofbauer, J., and S. Huttegger. 2008. "Feasibility of Communication in Binary Signaling Games." *Journal of Theoretical Biology* 254 (4): 843–49.
- Hu, Y., B. Skyrms, and P. Tarrès. 2011. "Reinforcement Learning in a Signaling Game." Unpublished manuscript, arXiv, Cornell University. <https://arxiv.org/abs/1103.5818>.
- Huttegger, S., B. Skyrms, P. Tarrès, and E. Wagner. 2014. "Some Dynamics of Signaling Games." *Proceedings of the National Academy of Sciences* 111 (S3): 10873–80.
- Lewis, D. 1969. *Convention*. Cambridge, MA: Harvard University Press.
- Manser, M., R. Seyfarth, and D. Cheney. 2002. "Suricate Alarm Calls Signal Predator Class and Urgency." *Trends in Cognitive Science* 6 (2): 55–57.
- Maynard Smith, J. 1965. "The Evolution of Alarm Calls." *American Naturalist* 99:59–63.
- Ouattaraa, K., A. Lemassona, and K. Zuberbühler. 2009. "Campbell's Monkeys Concatenate Vocalizations into Context-Specific Call Sequences." *Proceedings of the National Academy of Sciences of the USA* 106 (51): 22026–31.
- Roth, A. E., and I. Erev. 1995. "Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Immediate Term." *Games and Economic Behavior* 8:164–212.
- Searcy, W. A., and S. Nowicki. 2006. *The Evolution of Animal Communication: Reliability and Deception in Signaling Systems*. Princeton, NJ: Princeton University Press.
- Sherman, P. W. 1977. "Nepotism and the Evolution of Alarm Calls." *Science* 197:1246–53.
- Skyrms, B. 2006. "Signals." *Philosophy of Science* 75 (5): 489–500.
- . 2009. "Evolution of Signaling Systems with Multiple Senders and Receivers." *Philosophical Transactions of the Royal Society B* 364 (1518): 771–79.
- . 2010. *Signals: Evolution, Learning, and Information*. New York: Oxford University Press.
- Slobodchikoff, C. N., A. Paseka, and J. L. Verdolin. 2009. "Prairie Dog Alarm Calls Encode Labels about Predator Colors." *Animal Cognition* 12:435–39.
- Steinert-Threlkeld, Shane. 2016. "Compositional Signaling in a Complex World." *Journal of Logic, Language and Information* 25 (3–4): 379–97.
- Suzuki, Toshitaka N., David Wheatcroft, and Michael Griesser. 2016. "Experimental Evidence for Compositional Syntax in Bird Calls." *Nature Communications* 7, art. 10986. <https://doi.org/10.1038/ncomms10986>.
- Zehavi, A., and A. Zavahi. 1997. *The Handicap Principle: A Missing Piece of Darwin's Puzzle*. Oxford: Oxford University Press.