

Content in Simple Signalling Systems

Nicholas Shea, Peter Godfrey-Smith, and
Rosa Cao

ABSTRACT

Our understanding of communication and its evolution has advanced significantly through the study of simple models involving interacting senders and receivers of signals. Many theorists have thought that the resources of mathematical information theory are all that are needed to capture the meaning or content that is being communicated in these systems. However, the way theorists routinely talk about the models implicitly draws on a conception of content that is richer than bare informational content, especially in contexts where false content is important. This article shows that this concept can be made precise by defining a notion of functional content that captures the degree to which different states of the world are involved in stabilizing senders' and receivers' use of a signal at equilibrium. A series of case studies is used to contrast functional content with informational content, and to illustrate the explanatory role and limitations of this definition of functional content.

1 *Introduction*

2 *Modelling Framework*

3 *Two Kinds of Content*

3.1 *Informational content*

3.2 *Functional content*

4 *Cases*

4.1 *Case 1: Simplest case*

4.2 *Case 2: Partial pooling*

4.3 *Case 3: Bottleneck*

4.4 *Case 4: Partial common interest*

4.5 *Case 5: Deception*

4.6 *Case 6: A further problem arising from divergent interests*

5 *Discussion*

Appendix

1 Introduction

Recent years have seen dramatic advances in our understanding of communication and its evolution, through new models developed in biology, philosophy, linguistics, and economics. The models in these areas take different forms, but many can be seen as having a common theme. They show how sign-using interactions between senders and receivers are stabilized by means of selection processes that bear on sender and receiver behaviours.¹

Communication is usually thought to involve the production of signs or representations that have meaning, or content of some kind. Writers working in, or influenced by, the mathematical theory of information have sometimes wanted to set these issues aside, as irrelevant or positively unhelpful. Freeman Dyson claims that information theory's central dogma is that 'meaning is irrelevant' (Dyson [2011]; see also Shannon [1948], p. 379). Another recent discussion concurs:

When information theorists think about coding, they are not thinking about semantic properties. All of the semantic properties are stuffed into the codebook, the interface between source structure and channel structure, which to information theorists is as interesting as a phonebook is to sociologists. (Bergstrom and Rosvall [2011], p. 171)

In an important treatment of this topic, Skyrms ([2010]) argues that although questions of meaning and content are worth considering, a straightforward extension of basic ideas in information theory suffices to handle them. Signals have informational content when they change the probabilities of states of the world, or of a receiver's actions. Informational content exists whenever probabilities are changed in this way, regardless of what role the messages play; the informational content of a signal is represented by a vector that records, for each possible world state, how much the signal changes the probability of that state compared to its antecedent probability. This, for Skyrms, is all we need to recognize when thinking about content.

We agree that one way of understanding the content of signals in sender-receiver systems is by applying information-theoretic ideas in this way. But, we argue, there is also another approach to the interpretation of signals in systems of this kind, one tied to the way that actions guided by a signal have consequences that can stabilize signing behaviours.

Note first that whether signals have informational content, in Skyrms's sense, does not depend on whether they are part of a system with signs being used successfully to coordinate action with the state of the world. They would still carry informational content even if they were part of a

¹ (Lewis [1969]; Spence [1973]; Crawford and Sobel [1982]; Robson [1990]; Farrell and Rabin [1996]; Bergstrom and Lachmann [1998]; Maynard Smith and Harper [2003]; Huttegger *et al.* [2010]; Skyrms [2010]; Clark [2011]; Stegmann [2013]; Zollman *et al.* [2013]).

system in which the use of signals was not achieving anything useful at all, the system was far from equilibrium, and signals were giving rise to behaviours poorly matched with the world. Existing discussions in the modelling literature sometimes acknowledge, explicitly or tacitly, the appeal of a notion of content that is tied to the maintenance of equilibria in some way.

One response to this situation is to look for a view of content that combines informational and ‘functional’ considerations of this kind. This may well be fruitful, but our approach in this article is different. We will treat ‘informational content’ and ‘functional content’ as two separate and useful concepts, with distinct explanatory roles. Informational content involves probabilistic associations between signs and the world; functional content involves relations between signs and the world that figure in the stabilization of a system of sign use. The aim of the article is to analyse content in a way that is not guided by common-sense intuitions but by consideration of which notions of content are useful when thinking about signalling systems and their evolution.

The next section outlines the modelling framework used in the article. Subsequent sections describe the two kinds of content and then proceed through a series of cases that illustrate the two kinds of content and their complementary roles. The article aims to motivate a distinction between informational and functional content, but does not purport to be the last word on how functional content should best be formalized. In the discussion of some cases, we acknowledge some problems for our proposed formalization and provisionally sketch some ways it could be amended to overcome those limitations.

2 Modelling Framework

Our discussion is concerned with signalling systems that have the structure of a Lewis signalling game. David Lewis ([1969]) gave a model of signalling in which we assume two agents, a sender and a receiver, where the sender has access to information about the state of the world, but cannot act on it except to send signals of some kind. The receiver can see only the signals, but can act in a way that generates payoffs for both sides. The payoffs resulting from a receiver’s pairing of an act with a state of the world might be the same for sender and receiver or they might differ.

Lewis assumed that sender and receiver policies were rationally chosen in a situation of common knowledge. Brian Skyrms ([1996], [2010]) gave an evolutionary recasting of Lewis’s model. Rational choice was replaced by natural selection, or in some cases by simple forms of learning. Evolution, learning, and choice are all processes in which the consequences of behaviours can ‘feed back’ and re-shape the rules governing behaviour at later time-steps. The sender modifies (or maintains) its sender’s rule, which maps states of the world to signals; the receiver modifies (or maintains) its receiver’s rule,

which maps signals to acts. When a combination of a sender's and a receiver's rule is such that neither side can change their rule unilaterally and be better off, given what the other is doing, the system is in a Nash equilibrium. When a combination of rules is such that any unilateral change makes the changer worse off, the system is in a strict Nash equilibrium.

The Lewis–Skyrms model is related to models discussed in economics (Crawford and Sobel [1982]; Farrell and Rabin [1996]) and in evolutionary biology (Maynard Smith and Harper [1995], [2003]; Bergstrom and Lachmann [1998]; Zollman *et al.* [2013]). Models in economics have explored issues like honesty in advertising and the use of signals to help maintain cooperation (Spence [1973]; Robson [1990]). Honesty in signalling has also been a focus of biological and evolutionary models, investigating especially the way that a cost associated with a signal can enforce honesty. Both evolutionary and economic modelling have explored the consequences of divergence of interests between senders and receivers for the possibility and nature of signalling. Our discussion will be focused on the set-up described by Lewis and Skyrms, but many of our conclusions can be extended more broadly.

Formally, we are concerned with situations where there is an exogenously determined state of the world, $\{S_1, S_2, \dots\}$, a sender who can detect this state and has a range of signals or messages available, $\{M_1, M_2, \dots\}$, and a receiver who can see the signals and may use them when choosing among available actions, $\{A_1, A_2, \dots\}$. States of the world are associated with objective probabilities, $P(S_i)$. Combinations of acts and states are associated with payoffs for each agent, represented by matrices (introduced below in Section 4). A sender's rule is a mapping from states to messages; a receiver's rule is a mapping from messages to acts. Both these rules may be 'pure' or 'mixed'; a sender may, for example, respond to S_1 by always producing M_1 (a pure strategy), or perhaps by producing M_1 with probability p and M_2 with probability $1 - p$ (a mixed strategy). Our analysis of cases in this article will be simple. In general, we will note combinations of senders' and receivers' rules that are equilibrium states, states where neither side has any incentive to change their behaviour. In some cases, drawing on the work of others, we will give a richer description, which notes how a case behaves under some rule of evolutionary change. Much of our discussion is intended to be neutral, though, about the details of the selection process shaping the sender's and receiver's behaviours.

3 Two Kinds of Content

3.1 Informational content

An appealing way to think about the content of signals in sender–receiver systems is to draw on concepts from information theory (Shannon [1948];

Dretske [1981]). Signals carry information about states of the world when they change the probabilities of those states (Skyrms [2010]). The term ‘change’ here should not be understood as involving strange causal relations between signal and state, but merely the fact that the probability of a state conditional upon the signal is different from the unconditional probability of that state. A signal has content when it tells us something about how the world is, where ‘tells’ is a matter of changing probabilities, providing evidence. Dretske ([1981]) developed a view of this kind, but for a signal to have content, he required that it raise the probability of some state of the world to one. A signal says that the world is in S_2 , for example, if the probability of the world being in S_2 , given the signal, is one, and its probability independent of the signal is less than one. Skyrms ([2010]) outlines a more general view of the informational content of signals. A signal has informational content if it changes the probabilities of at least some states of the world, and its content is given by all the changes it makes to the probabilities of those states. So if a signal raises the probability of S_2 , but does not bring it to one, it can still tell us something about S_2 . For Skyrms, the kind of content where some states’ probabilities are reduced to zero is a special case (which he labels ‘propositional content’).

Skyrms adopts a particular format for representing the changes made by a signal to the probabilities of a set of states. For a set of states, $\{S_j\}$, the content of a signal, M_j , is constituted by the changes made to the probability of each state by the signal, where each ‘change’ is measured as the binary logarithm of the ratio of the conditional to the unconditional probability of that state. That is, the content of M_j is:

$$\langle \log_2(P(S_1|M_j)/P(S_1)), \log_2(P(S_2|M_j)/P(S_2)), \dots, \log_2(P(S_i|M_j)/P(S_i)), \dots \rangle .$$

So the content of a message is a vector. In the special case where a message reduces the probabilities of some states to zero, Skyrms labels those states in the vector with minus infinity. For example, if a message eliminates all but one of four initially equiprobable states, the content will be of the form $\langle -\infty, 2, -\infty, -\infty \rangle$. Then the content can be given in a familiar propositional form by disjoining the remaining states. Here the content of the signal is S_2 ; in another case, it might be S_2 -or- S_3 , and so on.

We follow Skyrms in thinking of content in general as given by a vector, with contents that definitively rule out some states being a special case, but we will do this with a simpler method than Skyrms’s. For us, the informational content of message M is the vector of post-signal probabilities of the states, $P(S_i|M)$. So in the case given above, where a message eliminates all but one of four initially equiprobable states, the content will be of the form $\langle 0, 1, 0, 0 \rangle$. Both Skyrms’s and our method have advantages and disadvantages (Godfrey-Smith [2012]). A disadvantage with using post-signal probabilities to represent

content is the fact that the content vector is well defined even if the message has not changed any probabilities, so $P(S_i) = P(S_i|M)$ for all i . Our response is to stipulate that in cases where all the states have their probabilities unchanged by a signal, the signal has no informational content. Our use of the posterior probability vector is motivated in part by the way it makes possible some formal comparisons between informational and functional content.

So the informational content of a signal is the distribution of probabilities of states of the world, conditional on that signal, with the proviso that at least some of these probabilities differ from the unconditional probabilities of the states. The informational properties of signals depend solely, then, on the unconditional probabilities of the states, together with the sender's rule. In cases where a message rules out some states of the world, a narrative summary of the content can be given (in the form ' S_1 ', or ' S_1 –or– S_2 '). When no states of the world are ruled out, a narrative summary would be vacuous. As Skyrms notes, a signal can carry information about both the states of the world perceived by the sender and about acts produced by the receiver. Here we will only discuss informational content about the state of the world.

3.2 Functional content

Signals have informational content (in both Skyrms's and our sense) whether a sender–receiver system is at equilibrium or not, and whether the signals are doing anything useful for the users or not. If a sender and receiver have rules configured so that the sender maps states to signals one-to-one, and the receiver maps signals to acts one-to-one, but in a way that guarantees that the act produced is the worst one possible in each state, signals have the same informational content they would have if the sender was performing the same mapping of states to signals, but the receiver was producing the best act in each state. Informational content is insensitive to facts about how well things are going and whether the system is at any kind of equilibrium.

This is not in any sense a problem for the notion of informational content. However, many writers have formed the view that content, of at least some variety, is dependent on those further factors. This might be seen as recognition of a richer concept of 'meaning' than mere informational content. For example, Simon Huttegger takes linguistic meaning ('the linguistic component of the truth of a statement') to be fixed by the conventions of meaning (Huttegger [2007a], p. 2), which are strict Nash equilibria of signalling games (p. 9).² Similarly, William Harms identifies 'primitive content' with

² Huttegger's ([2007b]) work on the distinction between indicative and imperative content also suggests that there is a role for functional considerations in defining content; see also (Zollman [2011]). These discussions may lead in the direction of an alternative notion of functional content to the one presented in the present article.

pairs of dispositions of senders to produce signals and receivers to act on signals, when such pairs have been stabilized by evolution or learning (Harms [2004]).³

These thoughts suggest that there is an additional way of thinking about content in signalling systems, having to do with the stabilization of the setup and the beneficial consequences of sender–receiver coordination. In the biological literature on animal signalling, the concept of ‘functional reference’ has been applied to such situations (Macedonia and Evans [1993]; Scarantino [2013]; cf. Wheeler and Fischer [2012]). In philosophy, both information-theoretic relationships and relationships involving success and stabilization of representation-using systems have been employed as the basis for general theories of content. They are usually seen as rivals: informational theories analyse content in terms of correlations between representations and states (Dretske [1981]; Fodor [1990]); teleosemantic theories hold that the content of a representation derives from the way a ‘consumer’ system acts on the representation to produce adaptive behaviour that has been relevant to the stabilization of that representation-using system (Millikan [1984], [1989]; Papineau [1984], [1993]). For example, when a vervet monkey sees a snake and makes a particular sound, ‘consumer’ monkeys run for cover in the trees. This has been useful in cases where the sound was produced in the presence of snakes, so ‘Snake!’ is the content of the sound, even if those cases are rare and many sounds are false alarms. Some philosophical theories of content rely on both functional and informational properties in combination (Neander [forthcoming]; Price [2001]; Shea [2007]).⁴

Those earlier debates about informational and teleofunctional theories were not generally carried out in the context of a sender–receiver model of the kind we are concerned with here.⁵ Rather than aiming for a choice between informational and functional properties, or a ‘gluing together’ of them, here we look at the idea that there are two kinds of content that messages can have in a sender–receiver system. One kind is derived from informational properties of the message—the way messages correlate with states of the world—and the other arises from the role the message plays in stabilization of the system through some process of selection.

Accordingly, we define functional content as follows: The messages in a sender–receiver system have functional content only if the system is at an

³ Indeed, Skyrms ([2010], p. 47) himself sometimes privileges the kind of information flow found at equilibrium—where the receiver ‘acts just as she would have if she had observed the state directly’ over other cases where just as much information is transmitted.

⁴ In the literature on ‘functional reference’ in animal communication, mentioned above, Scarantino ([2013], p. 1016) does the same in combining a ‘contextual perception criterion’ (dependent on evolutionary functions) with a ‘contextual information criterion’.

⁵ Harms ([2004]) was perhaps the first to connect sender–receiver models with a functional notion of content.

equilibrium maintained by some selection process.⁶ If it is, then for each signal M , we ask whether there is a behaviour (or distribution over behaviours) of the receiver specific to M , in the sense that the receiver responds differently to M than it does to some other available signal. (Note that this allows that the receiver may respond the same way to some other signal M' , but rules out that the receiver should respond the same way to all signals in the system.) If so, we look at whether there is a specific state of the world that obtains on some occasions when the message is sent, where the relation between that state of the world and the behaviours produced by the message contributes to the stabilization of those sender and receiver behaviours. If so, that state is the content of M . If the receiver's behaviour in response to M is stabilized by the obtaining of more than one world state on different occasions, the signal will have a disjunctive content involving all those world states.

In the case of informational content, we followed Skyrms in saying that content in general is given by a vector. We apply the same principle to functional content. The informational content vector takes the form of a list of entries that sum to one—the posterior probabilities of states of the world. The functional content vector we use here is also a list of entries that sum to one, though these entries are not probabilities. Whereas the informational content vector for a signal gives, for each state, how probable it is in the light of the signal, the functional content vector gives, for each state, the degree of involvement of that state in the stabilization of the sender's and receiver's behaviours regarding that signal. In the simplest cases, as with the vervet's 'Snake!' alarm call, there is just one state of the world whose obtaining figures in the stabilization of the system. But suppose that this particular alarm call has been mostly useful when there have been snakes around, and has afforded some protection when there are wild dogs around instead. Then the call has some functional involvement with both states.

More precisely, we define the functional content vector for a message in relation to baseline payoffs for the sender and receiver obtained in the absence of signalling. (The following recipe is expressed more formally in the appendix.) The baseline for each agent is the agent's average payoff in a situation where the receiver adopts the best strategy available to it without conditioning its behaviour on any signals (cf. Scott-Phillips *et al.* [2012], p. 1944). Non-zero entries in the vector for the functional content of a message correspond to states in which the message is sent and both agents receive above-baseline payoffs, given the receiver's rule for that message. For each such state, we calculate the difference between the sender payoffs received in that state and its baseline; we calculate the corresponding difference for the receiver. When

⁶ Birch ([2014]) uses a different way of defining content to argue that signals in out-of-equilibrium states have propositional content (which, as with our functional content, in general differs from informational content).

necessary, we take the smaller difference to yield a single value for each state. (See below for discussion of when this minimum must be considered and what role it plays.) These values are weighted by the posterior probabilities of the states, given the signal, and normalized to sum to one.⁷ The result is a vector representing the relative importance of each state to the stabilization of the sender and receiver rules for that message. This can be seen as a measure of the degree of involvement of the message with each state, given how the message is produced and used to guide action. (Some complications arise when sender and receiver payoffs differ, but do not differ so much that only one payoff is above baseline—we discuss these below.)

In that first presentation, we assumed that the receiver performs a single action in response to M . A receiver might ‘mix’ its behavioural responses to M , however, producing (say) act A_1 half the time and A_2 the rest of the time. In those cases, each action is analysed separately in the way outlined above, and the results are averaged, weighted by the probability that the receiver will produce the action in response to M .

The two kinds of content have the same form—distributions over states of the world, one reflecting posterior probabilities and one reflecting functional involvement. In cases where one or more entries are zero, a narrative summary of the content is available. This applies to both kinds of content. For example, a vector of the form $\langle 0, 0.6, 0.4 \rangle$ can be summarized as S_2 -or- S_3 . Vectors with no non-zero entries do not have a non-vacuous narrative summary.

Sometimes the informational content and functional content will coincide and, in some cases, will diverge. Lastly, the truth—the state of the world on an occasion when a signal is produced—can also be represented in the same form as the two kinds of content, with a distribution summing (trivially) to one. If, for example, there are three possible states of the world, S_1 , S_2 , and S_3 , and on some occasion S_2 is the actual state, this can be represented in a vector: $\langle 0, 1, 0 \rangle$. So the state of the world, the informational content of a signal, and the functional content of a signal all have the same form.

Before showing how these definitions play out in some cases from the existing literature, we will comment briefly on two alternative proposals. Harms ([2010]) illustrates a rather different way of connecting Lewis-style signalling games with philosophical work on teleosemantic theories of content. Harms does not use vectors to capture functional content. Our treatment also differs from Harms’s in making functional content partly a matter of the relative magnitude of the payoffs received in different states. Harms has a different focus, driven by concerns about how the world can be divided into states objectively. As a result, he dispenses with states of the world external to the sender–receiver system and characterizes his model only by reference to

⁷ This is equivalent to weighting the posterior probabilities by a function of the payoffs.

states of the sender's sensory apparatus and the payoffs that are received in those states. Functional contents are regions of a state space defined by the range of available sensory states and payoffs. There is not scope here to explore the extent to which Harms's approach is a rival to the one we develop here and the extent to which they are complementary.

Our functional content vector is broadly in the spirit of Oliver Lean's ([2014]) 'informational functions'. However, Lean casts his approach as contrasting with teleosemantic accounts of semantic information in biology, arguing that greater clarity is achieved by analysing function separately from information, and treating information in the style of Shannon. By contrast, we argue that a function-related notion of content is a useful resource for analysing communication in signalling systems.

We now turn to examples that illustrate the different roles of the two kinds of content.

4 Cases

4.1 Case 1: The simplest case

The simplest case is where there are two world states, two signals, and two acts, and the world states are equally probable. Both agents receive a positive payoff when A_1 is produced in S_1 and the same when A_2 is produced in S_2 , and neither receives a payoff otherwise. There are four possible sender strategies and four possible receiver strategies (leaving aside mixed strategies). Two of these are combinations of sender and receiver behaviours in which maximum payoff is achieved by both parties on every trial because signals are used to perfectly correlate the receiver's actions with the state of the world. One of these signalling systems is shown in Figure 1; here the sender invariably sends M_1 in response to S_1 , the receiver produces A_1 in response, and so on. The other simply swaps M_1 with M_2 in Figure 1. These are the only strict Nash equilibria of the game. Recent models have also shown that evolutionary processes can guide populations of various kinds to these equilibrium states (Huttegger *et al.* [2010]; Skyrms [2010]).




Figure 1. A signalling system in the case where there are two world states, two acts, and two signals available.

Table 1. Relations between informational and functional content for Case 1

		Informational content	Functional content
Messages	M_1	$\langle 1, 0 \rangle; S_1$	$\langle 1, 0 \rangle; S_1$
	M_2	$\langle 0, 1 \rangle; S_2$	$\langle 0, 1 \rangle; S_2$

Contents are given first in vector form and then in a narrative summary.

At the equilibrium shown in Figure 1, signal M_1 makes state S_1 certain and completely rules out state S_2 , so the post-signal probabilities are $\langle 1, 0 \rangle$. The functional content of M_1 is determined, as explained above, by examining the behaviour of the receiver specific to that signal and noting which pairing of messages to states contributes to the stabilization of the system. In this case, the functional contents of both messages are the same as their informational contents; M_1 is produced always and only in S_1 , and M_1 gives rise to A_1 , which contributes to the stabilization of the system if and only if S_1 obtains. So the contents are as set out in Table 1.

4.2 Case 2: Partial pooling

Even in simple situations like the set-up above, as soon as the probabilities of the two world states differ, informative signalling may become evolutionarily unlikely. In a pooling equilibrium, the sender sends the same signal in both states and so the signal is completely uninformative about the state of the world. Correspondingly, the receiver ignores the signal and performs the same action regardless. These equilibria exist even when the probabilities of states are equal, but they are more evolutionarily relevant when those probabilities are unequal, because in evolutionary models of situations in which the probabilities are unequal, populations do frequently end up in pooling equilibria. These are models in which each agent in the population plays the sender role half the time and the receiver role half the time, receiving payoffs according to the matching of receiver actions with states, and the population evolves by the replicator dynamics (Huttegger *et al.* [2010]; Skyrms [2010]). Pooling is a common outcome because agents implementing pairs of behaviours that constitute a signalling system incur a cost when they encounter pooling agents, since they then condition their behaviour on a completely uninformative signal. Simply performing the behaviour best suited to the most probable state is sufficiently profitable that it may be hard for signalling to invade.

Suppose we have a case like this, with S_1 much more probable than S_2 , where the sender sends M_1 in every state and the receiver performs A_1 , regardless of what they see. Then the signals do not change the probabilities of states of the world at all, in which case no signal has informational content in our sense.




Figure 2. A case of partial pooling in a system with three states, signals, and acts.

As the receiver performs the same acts in response to all messages, no signal has functional content either. As described above, a signal only has functional content when there is a characteristic behaviour resulting from that signal that plays a role in the stabilization of the system. Here, no signals are associated with characteristic behaviours in this sense.

Once there are three states, signals, and acts, partial pooling becomes possible, where the sender pools two world states together under the same signal, but sends a different signal in the third state.⁸ In the strategies shown in Figure 2, a case drawn from (Skyrms [2010]), the sender sends M_1 in response to both S_1 and S_2 , and mixes M_2 and M_3 in response to S_3 , with probabilities x and $1 - x$, respectively. The receiver maps both M_2 and M_3 to act A_3 , and mixes its response to M_1 , producing A_1 and A_2 with probabilities y and $1 - y$, respectively.⁹ Here we assume again that the three states of the world are equally probable. The assumptions about payoffs are as they were above: both actors receive a payoff in world state S_i if and only if act A_i is produced, with the magnitude of the payoffs the same in each case. In evolutionary simulations of the kind described above, some populations of senders and receivers do end up at equilibria of this kind (Huttegger *et al.* [2010]; Skyrms [2010]).

In this case, message M_1 shifts the probabilities equally towards both S_1 and S_2 , and M_2 and M_3 both shift the probabilities towards S_3 , giving rise to the informational contents set out in Table 2.

The receiver's behaviour in response to M_1 is to perform a mixture of A_1 and A_2 . What is the functional content of M_1 ? What is the condition whose obtaining on occasions where M_1 is acted on explains the success of this mixed policy of behaviour? The answer is that this depends on the value of y . In some situations, both S_1 and S_2 are involved in generating payoffs that are above baseline, given the mix of actions performed in response to M_1 . In those cases,

⁸ (Barrett [2006]) is the first discussion of pooling equilibria for signalling games such that, for any number, n , there are n states, n signals, and n acts; see also (Barrett [2007]).

⁹ In the cases we are considering, x and y are non-zero.

Table 2. Relations between informational and functional content for Case 2

		Informational content	Functional content
Messages	M_1	$\langle 0.5, 0.5, 0 \rangle; S_1\text{-or-}S_2$	if $y \geq \frac{2}{3}$, $\langle 1, 0, 0 \rangle; S_1$ if $\frac{2}{3} > y > \frac{1}{3}$, $\langle 3y - 1, 2 - 3y, 0 \rangle; S_1\text{-or-}S_2$ if $y \leq \frac{1}{3}$, $\langle 0, 1, 0 \rangle; S_2$
	M_2	$\langle 0, 0, 1 \rangle; S_3$	$\langle 0, 0, 1 \rangle; S_3$
	M_3	$\langle 0, 0, 1 \rangle; S_3$	$\langle 0, 0, 1 \rangle; S_3$

Contents are given first in vector form and then in a narrative summary.

the condition is disjunctive: the functional content of M_1 is $S_1\text{-or-}S_2$. This is a rough narrative summary, though. The functional content vector for M_1 is more specific, as it reflects the fact that proportion $3y - 1$ of the payoffs received at equilibrium are in world state S_1 and $2 - 3y$ in S_2 .

In other situations, when y is close to an extreme value, one or other of S_1 and S_2 does not play such a role, and the functional content is not disjunctive. The contents of the three signals are set out in Table 2.

As shown in Table 2, this case features divergence between functional and informational content, where the degree of divergence depends on y . When expressed in narrative terms, the functional content is stronger, for high and low values of y .

4.3 Case 3: Bottleneck

We now consider a different situation in which sender and receiver payoffs are suboptimal but the system can be at equilibrium. This is a case where there are not enough messages available to cover all the states—there are three world states but only two signals available by which to communicate about them. In the solution in Figure 3, action A_2 is never performed, and in S_2 , the agents receive the suboptimal reward of four, obtainable by performing A_1 in S_2 . This combination of behaviours produces the best outcome possible in the situation (Skyrms [2010], p. 113).

In this and subsequent cases, the details of payoffs are important. We represent them in a table with entries for the payoff received for each action in each world state. In Table 3, sender and receiver payoffs do not differ from one another.

The strategy in Figure 3 is structurally similar to the strategy in the previous case (Figure 2) in that it pools two world states. This is another case in which the functional content of M_1 differs from its informational content. Though




Figure 3. Sender and receiver behaviours in a ‘bottleneck’ case, with fewer messages than states.

Table 3. Payoffs for a ‘bottleneck’ case

		Acts		
		A_1	A_2	A_3
States	S_1	7	0	2
	S_2	4	6	0
	S_3	0	5	10

Table 4. Relations between informational and functional content for Case 3

		Informational content	Functional content
Messages	M_1	$\langle 0.5, 0.5, 0 \rangle$; S_1 -or- S_2	$\langle 1, 0, 0 \rangle$; S_1
	M_2	$\langle 0, 0, 1 \rangle$; S_3	$\langle 0, 0, 1 \rangle$; S_3

Contents are given first in vector form and then in a narrative summary.

the behaviour produced in response to M_1 does yield some payoff in S_2 , this payoff does not exceed the baseline achievable in the absence of signalling. For M_2 , in contrast, the functional and informational contents line up entirely (Table 4).

4.4 Case 4: Partial common interest

We now consider a game in which the payoffs for sender and receiver differ such that their interests are not fully aligned. They agree about the best action in one of the states (S_3), but in the other two, they have a different preference

Table 5. Payoffs in a case of partial common interest

		Acts		
		A_1	A_2	A_3
States	S_1	2, 10	0, 0	10, 8
	S_2	0, 0	2, 10	10, 8
	S_3	0, 0	10, 10	0, 0

Payoffs in each cell are to sender and receiver, respectively.




Figure 4. A case of partial common interest.

order.¹⁰ The payoffs are shown in Table 5.¹¹ In Figure 4, a combination of sender and receiver rules is shown that yields an equilibrium for this system (Skyrms [2010], p. 80). The sender uses M_1 to rule out S_3 and raise the probability of S_1 and S_2 equally, inducing the receiver to perform the sender’s preferred action in both states, since that action also pays off reasonably well for the receiver. The sender has an incentive not to differentiate S_1 and S_2 because then the receiver would perform its preferred action for each, to the detriment of the sender. So here imperfect alignment of interests produces partial pooling of states by the sender.

One reason researchers have been interested in cases where the interests of senders and receiver differ is because it raises the possibility of deception. However deception might be analysed in detail, at least the paradigm cases involve the sender using signals to achieve payoffs that run counter to the best interests of the receiver by inducing the receiver to perform actions that are not

¹⁰ For a discussion of measures of the degree of common interest based on divergence between the sender’s and receiver’s preference orderings over actions in states, see (Godfrey-Smith and Martínez [2013]).

¹¹ In the previous payoff matrix (Table 3), there was one entry for both sender and receiver payoffs in a combination of act and state. In Table 5, each cell contains a pair of numbers, for sender and receiver payoffs, respectively, for the corresponding act and state.

Table 6. Relations between informational and functional content for Case 4

		Informational content	Functional content
Messages	M_1	$\langle 0.5, 0.5, 0 \rangle; S_1\text{-or-}S_2$	$\langle 0.5, 0.5, 0 \rangle; S_1\text{-or-}S_2$
	M_2	$\langle 0, 0, 1 \rangle; S_3$	$\langle 0, 0, 1 \rangle; S_3$

Contents are given first in vector form and then in a narrative summary.

well aligned, given their interests, with the state of the world. Skyrms argues that the equilibrium shown in Figure 4 is a case of deception. We do not agree. What is true in this case is that signal M_1 carries less than perfect information about the actual state, failing to distinguish S_1 from S_2 . The receiver produces a cover-all behaviour that generates reasonably good payoffs in both S_1 and S_2 . In no circumstance does the receiver produce an action well suited only to one state when a different state obtains. The receiver's payoffs are always above their baseline. The functional and informational contents of the two messages used are given in Table 6. The sender is conveying and the receiver is acting on a true disjunctive content every time M_1 is sent ($S_1\text{-or-}S_2$).

Part of the reason Skyrms holds that this is a case of deception is the fact that when M_1 is sent, there is misinformation; the probability of a non-actual state of the world is raised by the signal (Skyrms [2010], p. 80). However, that was also true in the two cases of pooling discussed above (Cases 2 and 3), where signals, again, did not discriminate all states. We think this case (Case 4), is merely a case of strategic withholding of information by the sender, a phenomenon quite distinct from deception. We will next consider a case that we do regard as one of *bona fide* deception.

4.5 Case 5: Deception

To illustrate the possibility of genuine deception, we consider a signalling game relevant to animal communication, modifying a game discussed in (Zollman *et al.* [2013], Figures 2 and 3, Table 2). Suppose senders are males and receivers are females, and males signal to advertise their quality. Males can be high or low quality. Males always prefer to mate, whereas females prefer to mate only with high-quality males. (These contexts involving display are assumed to not be the only contexts in which females can mate; uniform refusal to mate by a female in these contexts does not imply zero fitness.) Suppose too that males have a signal available that is more costly for low-quality than high-quality individuals to send. (The payoffs are represented in Table 7.) Then a stable signalling system can evolve in which males reliably signal their quality and females condition their mating behaviour on the signal.

Table 7. Payoffs in the deception case described in Section 4.5

		Acts	
		A_1 (mate)	A_2 (not mate)
States	S_1 (high-quality male)	2, 2	1, 1
	S_2 (low-quality male)	2, 0	1, 1

Payoffs in each cell are to sender (male) and receiver (female), respectively.




Figure 5. A case of deception: a hybrid equilibrium of the case described in Section 4.5. S_1 and S_2 are the possible states of the male sender. M_1 is a costly signal. It costs $\frac{1}{2}$ for low-quality males (S_2) to send M_1 , but only $\frac{1}{4}$ for high-quality males (S_1). Signal M_2 has no cost.

There is also a ‘hybrid equilibrium’ of this game, which we will focus on here, in which both senders and receivers sometimes mix their behaviours and sometimes do not. High-quality male senders always send the more costly high-quality signal. Low-quality males randomize, sending the high-cost signal in some cases and the low-cost low-quality signal on other occasions. On the receiver side, males who send the low-quality signal are always rejected and those who send the high-quality signal are accepted with some probability and rejected the rest of the time. Whether a hybrid equilibrium exists depends on the parameter values—payoffs, costs of signals, and the frequency of high-quality males—and this equilibrium will involve a specific mix of sender and receiver behaviours. One example of a set of parameters for which an equilibrium exists is given in Table 7 and Figure 5. Here, the low-quality males send the high-cost signal with probability $\frac{2}{3}$ and high-cost signals are accepted with probability $\frac{1}{2}$. This combination of sender and receiver strategies is a Nash equilibrium.¹²

¹² The equilibrium requires that the probability of the receiver accepting the high-cost signal (that is, performing A_1 in response to M_1) is equal to the cost to low-quality senders of sending the high-cost signal (that is, of sending M_1 in S_2). Both are equal to $\frac{1}{2}$ in our illustration. This has the effect of ensuring that the benefit to low-quality individuals of sometimes achieving a mating

Table 8. Relations between informational and functional content for Case 5

		Informational content	Functional content
Messages	M_1	$\langle 0.5, 0.5 \rangle$; no propositional content	$\langle 1, 0 \rangle$; S_1
	M_2	$\langle 0, 1 \rangle$; S_2	None

Contents are given first in vector form and then where possible in a narrative summary.

The informational and functional contents of messages at this hybrid equilibrium are set out in Table 8. Message M_1 has no propositional informational content, because no state is ruled out by the message. However, it does have a functional content that is propositional: S_1 . This is the only state in which the receiver’s rule at equilibrium generates for both sides an above-baseline payoff. As a consequence, when M_1 is sent in S_2 , which does happen some of the time, this message has false propositional content. It says the world is in S_1 when in fact the world is in S_2 . False propositional content is quite different from what Skyrms called ‘misinformation’. The case in Figure 5 is the only case so far in which a signal sometimes has false propositional content, while misinformation in Skyrms’s sense is found also in cases with bottlenecks and pooling (Sections 4.2–4.4).

We understand deception to occur when a message with a false content is sent and the receiver is induced to behave in a way that benefits the sender and harms the receiver. ‘Deception’ in this sense is a success term; it can be distinguished from attempted deception, which occurs when a message with false content is sent in a way that has the potential to benefit the sender at the expense of the receiver. So, for example, when the sender sends M_1 in S_2 but the receiver refuses to mate, this is merely a case of attempted deception. If the receiver does mate with the low-quality sender, this is a case of deception.

Existing discussions of cases of this kind routinely assume a concept of deception similar to ours, without spelling out a view of content that licences it. For example, Zollman *et al.* ([2013]) describe the hybrid equilibria that can exist in these signalling games in the following terms: ‘In plain English, this means that the sender sometimes “lies” and is honest at other times, whereas the receiver only sometimes chooses the sender’s favoured action’. If the notion of ‘lying’ requires that a message has false content, and not merely that it withholds some information, then informational content as discussed here and elsewhere does not suffice to make sense of lying, and something like functional content in our sense is needed.

is exactly balanced, on average, by the cost to low-quality individuals of sending the high-cost signal.

A further notable feature of our treatment of this case is that the functional content of M_2 is undefined, as Table 8 shows. This is because no state of the world generates higher-than baseline payoffs given the receiver's equilibrium response to M_2 . Indeed, although M_2 is treated in the model as a signal, it is associated neither with costs nor the possibility of benefit, so it is more naturally understood as the absence of a signal—as a 'null' signalling behaviour.

While that is a satisfactory result in the present case, in other games with intrinsic signalling costs, our proposed definition of functional content is more problematic. Bergstrom and Lachmann ([1997]) analyse another game with costly signals, the Sir Philip Sidney game. They show that there are separating equilibria in which both sender and receiver are worse off than they would be if the receiver produces its best cover-all response to completely uninformative signals.¹³ In such a separating equilibrium, there is no state in which both players obtain payoffs above their baseline, as we have defined the baseline; so there is no functional content.

To sketch a response to this problem, we return to the theoretical motivation for our account. Functional content is a matter of more than just coordinating actions with the state of the world. That happens in the case of perfect anti-signalling mentioned above, where signals are perfectly coordinated with world states, but no payoffs result. Functional contents arise where the players coordinate actions with the state of the world successfully. Isolating cases of successful coordination calls for a standard of comparison, which is what our baselines achieve. If one accepts this theoretical motivation, then it follows that functional content is not ubiquitous—it is absent in some equilibria where signals are coordinated with the state of the world.

As formulated, our definition has the consequence that functional content is absent when players fall into an equilibrium in a costly signalling game that makes them worse off in every world state than they would be without signalling (although there is functional content in the costly signalling game we analyse here). Rather than just accepting that consequence, another solution would be to define baselines more locally when there are intrinsic signalling costs, in terms of nearby states in which both sides do worse than at the equilibrium. We do not attempt to resolve this issue here.

4.6 Case 6: A further problem arising from divergent interests

When sender and receiver interests diverge, but do not diverge greatly, a problem can arise that has not been addressed in our cases above. This problem comes when, given some act or mix of acts produced in response to a message at equilibrium, sender and receiver both achieve above-baseline payoffs in the

¹³ We are grateful to a referee for pointing out the implications of this game.

Table 9. Payoffs in Case 6

		Acts				
		A_1	A_2	A_3	A_4	A_5
States	S_1	2, 2	5, 0	5, 0	0, 0	0, 0
	S_2	2, 2	0, 5	0, 5	0, 0	0, 0
	S_3	2, 2	0, 0	0, 0	5, 0	5, 0
	S_4	2, 2	0, 0	0, 0	0, 5	0, 5

Payoffs in each cell are to sender and receiver, respectively.

same combination of states, but the degrees to which they benefit in each of these states differ. Then when a vector representation of functional content is given, strictly speaking there will be one vector for the sender and one for the receiver, not a single vector describing both. This does not happen in either of the two cases with divergent interests discussed above. In one of these cases (Case 5), no message is interpreted in a way that gives both parties an above-baseline payoff in more than one state (only the sender receives an above-baseline payoff in more than one state, given the receiver's rule for M_1). In the other Case 4, both sender and receiver obtain above-baseline payoffs in S_1 and S_2 , given the rules associated with M_1 , and these payoffs do differ between sender and receiver, but for neither agent is one state preferable to the other. So there is no qualitative difference between the agents with respect to the roles of S_1 and S_2 in stabilizing this aspect of their interaction.

In other possible cases, when the interests of the agents diverge in a way that leads to a message being associated with more than one state of the world for each agent, though with different weightings for these states across the two agents, the formulation we give is designed to capture the 'overlap' between sender and receiver interests (see the appendix for details). As we noted, in such cases it is also straightforward to record separate functional content vectors for sender and receiver, respectively. Comparisons between our preferred functional content vector, which captures the overlap, and the separate functional content vectors for sender and receiver would show the respects in which sender and receiver have different interests in the way the signal is connected to world states at equilibrium.

A case put forward by a referee helpfully illustrates another way this kind of divergence can arise. In this new case, there are four equiprobable states, five available acts, and two costless signals, with payoffs as given in Table 9. We focus on the equilibrium shown in Figure 6, in which both players receive an above-baseline payoff in one particular state for each signal, but the relevant state differs for each player. For example, when M_1 is sent, the sender receives a payoff only in S_1 and the receiver only in S_2 . Our 'overlap' functional




Figure 6. Sender and receiver behaviours in the equilibrium considered in Case 6.

Table 10. Relations between informational and functional content for Case 6

	Informational content	Functional content for sender	Functional content for receiver
Messages	M_1 $\langle 0.5, 0.5, 0, 0 \rangle$; S_1 -or- S_2 M_2 $\langle 0, 0, 0.5, 0.5 \rangle$; S_3 -or- S_4	$\langle 1, 0, 0, 0 \rangle$; S_1 $\langle 0, 0, 1, 0 \rangle$; S_3	$\langle 0, 1, 0, 0 \rangle$; S_2 $\langle 0, 0, 0, 1 \rangle$; S_4

Contents are given first in vector form and then in a narrative summary.

content vector is undefined. Table 10 records separate functional content vectors for sender and receiver.

Comparing the separate functional content vectors for sender and receiver, and in the absence of an ‘overlap’ functional content vector, we can see that the two players have completely different interests in the way the signal is connected with world states at equilibrium. The receiver is only interested in the way M_1 carries information about state S_2 , whereas the sender receives a payoff only when S_1 obtains. An alternative perspective on this case would be to argue that sender and receiver do share an interest when M_1 is sent—an interest in the fact that S_1 -or- S_2 obtains. A natural move here would be to describe the game in a more coarse-grained way, so that S_1 -or- S_2 counts as a single state. Sender and receiver would then overlap in functional content with respect to that state. The difficulty is to formulate a rule for when it is appropriate to move to a more coarse-grained functional content vector that does not have the result that all cases of partial pooling turn into cases of perfect signalling with more coarse-grained states. While this case is clearly another reason to distinguish functional contents for sender and receiver in some cases, we have no settled view as to whether there is also a principled way to define a non-vacuous overlap functional content vector in this case.

5 Discussion

The small selection of examples above show that there is an important role for a notion of content that goes beyond purely informational content, even in these simple cases. Specifically, there is a role for a treatment that is connected to equilibria and how they are stabilized by payoffs. The way theorists routinely talk about simple signalling systems makes this clear. They say things that implicitly draw on a richer notion of content than informational content. This might be seen as metaphorical. But we have shown that a concept like this can be made precise and shown to be useful, especially in contexts where false content is important, such as in the analysis of deception.

Teleosemantics also aimed to capture the involvement signs have with the world. The concept of functional content developed here is a fine-grained take on that idea. The need to go beyond a purely informational treatment and introduce a broadly functional notion of content is one of the insights of Millikan ([1984]), Papineau ([1993]), and Dretske ([1988]). What we're doing is combining those ideas with Skyrms's introduction of a fine-grained vector representation of content. Our functional content vector captures the relative importance of different states when more than one state is involved in stabilizing a pattern of sender and receiver behaviours.

The concept of functional content we have developed here is not the only way this could be done. And it is clear that our treatment in this article still faces some problems. We hope to have shown that it is widely applicable enough to illustrate that there is space for an account of functional content alongside that of informational content.

Lastly, we make a comment about the status of these properties, which we have been calling a kind of 'content'. Clearly the signs themselves and their associated behaviours are much simpler and more rudimentary than those associated with human language and thought. They are probably simpler than most non-human sign systems as well. We don't claim that informational and functional content exhaust the rich semantic properties seen in language and thought. They can be thought of as simpler members of a family of semantic properties, or as precursors to real semantic properties. These simpler semantic or proto-semantic properties are, however, important features of signalling systems. Our notion of functional content captures a theoretically important aspect of sender–receiver interaction.

Appendix

Here we will provide the definition of the functional content vector.

$$\text{Functional content of signal } M = \left\langle \frac{x_1}{\sigma}, \frac{x_2}{\sigma}, \dots, \frac{x_n}{\sigma} \right\rangle.$$

We define the functional content vector for an arbitrary signal M , following the procedure given in Section 3.2 above. The vectors listed in the case studies above are the result of applying this procedure to each signal M_1, M_2, \dots found in the model. Below we proceed in two parts. First we define the baseline payoff for a signalling game. The baseline is then used as a threshold to generate components of the functional content vector.

A.1 Baseline Payoffs

We define the functional content vector in relation to the baseline payoffs obtained for the sender and receiver in the absence of signalling. Baseline for each (\bar{v}_r, \bar{v}_s) is its expected payoff given A^* , the action dictated by the best strategy the receiver can adopt without conditioning its behaviour on any signals.¹⁴ In defining the baseline here, we consider only pure receiver strategies since in the absence of signals, the receiver can never do better by mixing than by pursuing some pure strategy.

$$\text{Receiver's baseline payoff, } \bar{v}_r = \sum_{i=1}^n P(S_i) v_r(A^*|S_i),$$

$$\text{Sender's baseline payoff, } \bar{v}_s = \sum_{i=1}^n P(S_i) v_s(A^*|S_i),$$

$v_r(A^*|S_i)$ = Receiver's payoff for action A^* when world is in state S_i ,

$v_s(A^*|S_i)$ = Sender's payoff for action A^* when world is in state S_i ,

n = number of world states S_i .

A.2 Functional Content Vector

Components x_i in the functional content vector reflect the average payoff received from world state S_j when signal M is sent, thresholded by reference to the baseline payoffs calculated above. Non-zero entries correspond to states in which both agents receive above-baseline payoffs given the receiver's rule for M , and record the amount by which the threshold is exceeded. The requirement that both agents receive above-baseline payoffs implies that the agents have similar payoff matrices to some degree, but differences are still possible. Accordingly, we construct the vector entries by using the lesser of the amounts by which the two agents' payoffs surpass their baseline: $d_{\min}(A_j|S_i)$.

¹⁴ If there is more than one such action, then A^* can be any one of the receiver's best cover-all strategies, unless there are differences in the average payoff that results for the sender. We don't here attempt to resolve the question of how baselines should be defined in the latter case.

This is designed to represent the overlap between the sender’s and receiver’s interests:

$$d_{\min}(A_j|S_i) = \min((v_r(A_j|S_i) - \bar{v}_r), (v_s(A_j|S_i) - \bar{v}_s)).$$

This formulation is appropriate when there is a single population of agents that play the sender role half the time and the receiver role half the time. If there are separate populations of senders and receivers, then the payoffs in the payoff matrix should be transformed to a common scale. Before calculating the baselines and d_{\min} , the sender’s payoffs should be linearly transformed so that its maximum payoff is one, and the same for the receiver. For simplicity we do not include such transformations in the definition below.

The entries in the functional content vector are then defined as follows:

$$x_i = \begin{cases} P(S_i|M) \left\{ \sum_{j=1}^m P(A_j|M) d_{\min}(A_j|S_i) \right\}, \\ \text{if } \sum_{j=1}^m P(A_j|M) v_r(A_j|S_i) > \bar{v}_r \text{ and } \sum_{j=1}^m P(A_j|M) v_s(A_j|S_i) > \bar{v}_s \\ 0, \text{ otherwise} \end{cases}$$

m = number of available actions A_j

Finally, the components are normalized by σ so that, as with the informational content vector, they sum to one:

$$\sigma = \sum_{i=1}^n x_i.$$

The use of d_{\min} terms is only essential in special cases; see Section 4.6 of the main text for discussion. In such cases, another option would be to define separate functional content vectors for sender and receiver when they differ. We do not hold that one of these approaches is better than the other; each might represent different features of these cases.

Acknowledgements

This work is fully collaborative and all authors contributed equally. For comments on aspects of this material, the authors would like to thank audiences at the University of Aberdeen, the ‘Signalling and Meaning’ workshop at CUNY Graduate Center, UC Irvine, the British Society for the Philosophy of Science, the European Society for Philosophy and Psychology, and King’s College London, as well as Manolo Martinez, Ron Planer, and two anonymous

referees. This research was kindly supported by the Arts and Humanities Research Council (grant no. AH/M005933/1 to Nicholas Shea).

Nicholas Shea
Institute of Philosophy, School of Advanced Study
University of London
London, UK
nicholas.shea@sas.ac.uk

Peter Godfrey-Smith
Graduate Center
City University of New York
New York, USA
and
Unit for the History and Philosophy of Science
University of Sydney
Sydney, Australia
peter.godfrey-smith@sydney.edu.au

Rosa Cao
Philosophy Department
Stanford University
Stanford, USA
rosacao@stanford.edu

References

- Barrett, J. A. [2006]: ‘Numerical Simulations of the Lewis Signalling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar’ available at <escholarship.org/uc/item/5xr0b0vp>.
- Barrett, J. A. [2007]: ‘Dynamic Partitioning and the Conventionality of Kinds’, *Philosophy of Science*, **74**, pp. 527–46.
- Bergstrom, C. T. and Lachmann, M. [1997]: ‘Signaling among Relatives, I: Is Costly Signalling Too Costly?’, *Philosophical Transactions of the Royal Society B*, **352**, pp. 609–17.
- Bergstrom, C. T. and Lachmann, M. [1998]: ‘Signaling among Relatives, III: Talk Is Cheap’, *Proceedings of the National Academy of Sciences United States of America*, **95**, pp. 5100–5.
- Bergstrom, C. T. and Rosvall, M. [2011]: ‘The Transmission Sense of Information’, *Biology and Philosophy*, **26**, pp. 159–76.
- Birch, J. [2014]: ‘Propositional Content in Signalling Systems’, *Philosophical Studies*, **171**, pp. 493–512.
- Clark, R. [2011]: *Meaningful Games*, Cambridge, MA: MIT Press.
- Crawford, V. P. and Sobel, J. [1982]: ‘Strategic Information Transmission’, *Econometrica*, **50**, pp. 1431–51.
- Dretske, F. [1981]: *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.

- Dretske, F. [1988]: *Explaining Behavior: Reasons in a World of Causes*, Cambridge, MA: MIT Press.
- Dyson, F. [2011]: 'How We Know', *New York Review of Books*, **58**, p. 74.
- Farrell, J. and Rabin, M. [1996]: 'Cheap Talk', *The Journal of Economic Perspectives*, **10**, pp. 103–18.
- Fodor, J. A. [1990]: 'A Theory of Content II: The Theory', in his *A Theory of Content and Other Essays*, Cambridge, MA: MIT Press.
- Godfrey-Smith, P. [2012]: 'Review: *Signals: Evolution, Learning, and Information*, by Brian Skyrms', *Mind*, **120**, pp. 1288–97.
- Godfrey-Smith, P. and Martínez, M. [2013]: 'Communication and Common Interest', *Plos Computational Biology*, **9**, p. e1003282, available at <dx.Doi.Org/10.1371/Journal.Pcbi.1003282>.
- Harms, W. F. [2004]: 'Primitive Content, Translation, and the Emergence of Meaning in Animal Communication', in D. Kimbrough and U. Griebel (eds), *Evolution of Communication Systems: A Comparative Approach*, Cambridge, MA: MIT Press, pp. 31–48.
- Harms, W. F. [2010]: 'Determining Truth Conditions in Signalling Games', *Philosophical Studies*, **147**, pp. 23–35.
- Huttegger, S. M. [2007a]: 'Evolution and the Explanation of Meaning', *Philosophy of Science*, **74**, pp. 1–27.
- Huttegger, S. M. [2007b]: 'Evolutionary Explanations of Indicatives and Imperatives', *Erkenntnis*, **66**, pp. 409–36.
- Huttegger, S. M., Skyrms, B., Smead, R. and Zollman, K. J. S. [2010]: 'Evolutionary Dynamics of Lewis Signalling Games: Signalling Systems vs. Partial Pooling', *Synthese*, **172**, pp. 177–91.
- Lean, O. M. [2014]: 'Getting the Most out of Shannon Information', *Biology and Philosophy*, **29**, pp. 395–413.
- Lewis, D. [1969]: *Convention*, Cambridge, MA: Harvard University Press.
- Macedonia, J. M. and Evans, C. S. [1993]: 'Essay on Contemporary Issues in Ethology: Variation among Mammalian Alarm Call Systems and the Problem of Meaning in Animal Signals', *Ethology*, **93**, pp. 177–97.
- Maynard Smith, J. and Harper, D. [1995]: 'Animal Signals: Models and Terminology', *Journal of Theoretical Biology*, **177**, pp. 305–11.
- Maynard Smith, J. and Harper, D. [2003]: *Animal Signals*, Oxford: Oxford University Press.
- Millikan, R. G. [1984]: *Language, Thought, and Other Biological Categories*, Cambridge, MA: MIT Press.
- Millikan, R. G. [1989]: 'Biosemantics', *Journal of Philosophy*, **86**, pp. 281–97.
- Neander, K. [forthcoming]: *Mental Representation*, Cambridge, MA: MIT Press.
- Papineau, D. [1984]: 'Representation and Explanation', *Philosophy of Science*, **51**, pp. 550–72.
- Papineau, D. [1993]: *Philosophical Naturalism*, Oxford: Blackwell.
- Price, C. [2001]: *Functions in Mind*, Oxford: Oxford University Press.
- Robson, A. J. [1990]: 'Efficiency in Evolutionary Games: Darwin, Nash, and the Secret Handshake', *Journal of Theoretical Biology*, **144**, pp. 379–96.

- Scarantino, A. [2013]: 'Rethinking Functional Reference', *Philosophy of Science*, **80**, pp. 1006–18.
- Scott-Phillips, T. C., Blythe, R. A., Gardner, A. and West, S. A. [2012]: 'How Do Communication Systems Emerge?', *Proceedings of the Royal Society B*, **279**, pp. 1943–9.
- Shannon, C. E. [1948]: 'A Mathematical Theory of Communication', *Bell System Technical Journal*, **27**, pp. 379–423.
- Shea, N. [2007]: 'Consumers Need Information: Supplementing Teleosemantics with an Input Condition', *Philosophy and Phenomenological Research*, **75**, pp. 404–35.
- Skyrms, B. [1996]: *Evolution of the Social Contract*, Cambridge: Cambridge University Press.
- Skyrms, B. [2010]: *Signals: Evolution, Learning, and Information*, New York: Oxford University Press.
- Spence, M. [1973]: 'Job Market Signalling', *The Quarterly Journal of Economics*, **87**, pp. 355–74.
- Stegmann, U. [2013]: *Animal Communication Theory: Information and Influence*, Cambridge: Cambridge University Press.
- Wheeler, B. C. and Fischer, J. [2012]: 'Functionally Referential Signals: A Promising Paradigm Whose Time Has Passed', *Evolutionary Anthropology: Issues, News, and Reviews*, **21**, pp. 195–205.
- Zollman, K. J. S. [2011]: 'Separating Directives and Assertions Using Simple Signalling Games', *Journal of Philosophy*, **63**, pp. 158–69.
- Zollman, K. J. S., Bergstrom, C. T. and Huttegger, S. M. [2013]: 'Between Cheap and Costly Signals: The Evolution of Partially Honest Communication', *Proceedings of The Royal Society B*, **280**, 20121878.