



Signals without teleology

Carl T. Bergstrom^{a,b}, Simon M. Huttegger^c, Kevin J.S. Zollman^{d,*}

^a University of Washington, USA

^b Santa Fe Institute, USA

^c University of California, Irvine, USA

^d Carnegie Mellon University, USA

ABSTRACT

"Signals" are a conceptual apparatus in many scientific disciplines. Biologists inquire about the evolution of signals, economists talk about the signaling function of purchases and prices, and philosophers discuss the conditions under which signals acquire meaning. However, less attention has been paid to what *is* a signal. Most existing accounts are teleological in some way. This paper provides a definition of signals that avoids reference to form or purpose. Along the way we introduce novel notions of "information revealing" and "information concealing" moves in games. In the end, our account offers an alternative to teleological accounts of communication.

1. Introduction

"Clouds mean rain"

"The word 'dog' means dog"

Distinguishing between the two senses of the verb 'to mean' in the above sentences has occupied a central place in the philosophical discussion of language. The difference was captured by Peirce's distinction between indices — those signs which naturally correlate to the object they represent — and symbols — those signs that represent an object because of a certain arbitrary choice by a linguistic community (Peirce, 1867).

This distinction is carried forward in Grice's definition of non-natural meaning (Grice, 1957). Non-natural meaning is present when the speaker intends to communicate and intends that the listener recognize this intention, while natural meaning can exist without either of these. The second part of this account—that the speaker intends for the hearer to recognize her intention to communicate—is critical because it distinguishes between cases where the hearer merely observes an action that is performed for some purpose other than communicating. So while one might infer from my carrying an umbrella that it will rain, my carrying an umbrella is not the same as my telling you it will rain. In ordinary cases where I carry an umbrella, I do not intend to communicate anything nor do I care what you infer from it, whereas when I say to you directly: "It will rain," one supposes that I do.

Grice's notion of non-natural meaning is related to the distinction between conventional and non-conventional behaviors. Starting with (Lewis, 1969), conventional behaviors are defined (loosely) as those which are done exclusively because others do them. Language features many conventional choices: I say "It will rain" when I think it will rain only because other people also use those words in that way. If people used words differently, I would say something different. Carrying an umbrella, on the other hand, is not conventional: I do it because it keeps me dry irrespective of whether other people do it. Many conventional linguistic behaviors contain Gricean non-natural meaning.¹

A similar issue arises in biology. Many biologists are interested in distinguishing between signals—those phenotypes which have evolved solely to transmit information—and cues—those phenotypes that transmit information, but evolved for some other purpose (Maynard Smith & Harper, 2003).² Gricean intentions will not do the trick here, however. While primates might perhaps be said to have intentions, it is less obvious that bees do and highly likely that bacteria do not. Biologists would appreciate a distinction which accounts for all cases of the transmission of information between, and even within, organisms. Researchers who tackle this problem have, in one way or another, relied on a notion of biological function. Stated loosely, signals evolved because of their effectiveness in transmitting information while cues did not.

The biological way of capturing this distinction has been brought back into the philosophy in various guises. Dretske (Dretske, 1981) and

* Corresponding author.

E-mail address: zkollman@andrew.cmu.edu (K.J.S. Zollman).

¹ While often conventionality is equated with arbitrariness, it need not (Simons & Zollman, 2019). The term 'conventional' also has a complex history in biology (Guilford & Dawkins, 1995).

² Maynard Smith and Harper used an idiosyncratic definition of 'index' that is incompatible with the Peircean terminology. Here we stick to Peirce's definitions.

<https://doi.org/10.1016/j.shpsc.2020.101310>

Received 25 July 2018; Received in revised form 8 May 2020; Accepted 21 May 2020

1369-8486/ © 2020 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Skyrms (Skyrms, 2010) have used it in epistemology. In the philosophy of language, several authors have attempted to ground meaning in various amalgamations of information-theoretic concepts and notions of evolutionary history or biological function. In this later camp are the teleosemantic theories of Millikan (Millikan, 2005), Harms (Harms, 2004), Stegmann (Stegmann, 2009), Scott-Phillips (Scott-Phillips, 2008), Birch (Birch, 2014), and Shea, Godfrey-Smith, and Cao (Shea, Godfrey-Smith, & Cao, 2016). Under this description, one can view the teleosemantic enterprise as replacing Grice's intentions with biological or cultural function.³

Finally, the same set of problems emerges in the study of economic behavior. Since Spence's work (Spence, 1973) in the 1970's economists have come to recognize that many economic actions like setting prices or paying for celebrity endorsements might function to convey information from one party to another. Here, too, a critical distinction arises between those actions that are taken for the purpose of conveying information and those that convey information only incidentally. This distinction separates those economic behaviors which are truly signals from those that have another economic motivation. The study of "signaling games" has become a significant enterprise in economic thought. Recently, there is interest in so-called "self-signaling" where an economic action — like donating to a charity — is taken to purportedly send a signal to oneself at a future time (Gneezy, Gneezy, Riener, & Nelson, 2012).

All the methods for drawing the distinction create complexities for those interested in determining whether a particular sign is an index or a signal. One must determine the provenance of a sign by finding either the intentions that were present at its production or the evolutionary processes that resulted in its fixation. In some cases this may be possible, but in other situations it might be difficult. For instance, in studies of animal behavior scientists may not know the evolutionary history leading to the production of a sign, but at the same time they may lack the means to construct an account in terms of intentionality. Nevertheless, they often wish to talk meaningfully about an animal sending a signal or even deceiving another organism. In these difficult situations it would be very useful to have a method by which one could draw the distinction between signals and indices without relying on evolutionary history or speakers' intentions.

In the present paper we develop an alternative methodology for capturing the distinction between signals and signs, a methodology which relies solely on the game-theoretic structure of the interaction. Utilizing the canonical two-person models of communication that are studied in philosophy, biology, and economics, we define an "informational move." We say a move is informational when *you would not do it if the other party already knew everything you do*. By making this notion precise we will be able to distinguish between signals and indices without relying on teleology in one guise or another. Formalizing the notion of informational moves also allows us to uncover a communicative flip-side to the production of signals. In contrast to typical signals that serve to reveal information, we describe this new class of behaviors as "information concealing".

This represents two improvements on existing theories. First, in some cases it may be easier to determine whether or not something counts as a signal—namely, when the evolutionary or intentional provenance of a behavior may be empirically inaccessible. Second, it is simpler. Most teleosemantic theories require an understanding of the strategic situation (the game being played), as does our approach. But we do not require the kind of knowledge of evolutionary history that is

part and parcel of teleosemantics.

2. Signals

Peacocks vary in quality; some are desirable mates for peahens while others are not.⁴ Peahens want to mate with the good males and avoid the bad ones, whereas both high and low quality males would like to mate with a peahen. High quality males have an interest in finding a way to communicate their quality to peahens that cannot be imitated by their low quality rivals.

This example illustrates a general problem faced by human and non-human animals alike: receivers and some signalers would benefit from honest communication, but other signalers have incentives to be deceptive. How can honest communication be ensured? The canonical solution to this problem, called Costly Signaling Theory, was developed independently in biology and economics and employs the formal tools of game theory.⁵

Consider a simple model. Nature flips a (potentially biased) coin to determine the quality of the peacock: either high or low. The peacock learns his quality and takes one of two actions: he can either grow a big tail or grow a small tail. Growing a big tail requires substantial energetic expense, although the expense is considerably larger for the low quality peacock than it is for the high quality one. The peahen, who cannot directly observe the quality of the male, observes the size of the male's tail and decides whether to mate with him. The peahen has no intrinsic interest in the peacock's tail, only an interest in mating with the high quality male.

This story is summarized in the extensive form game pictured in Fig. 1. Here the benefit of mating for the male peacock is 1, but growing a large tail has a cost of c for the high quality male or d for the low quality male. For the female, mating with the high quality male pays 1 while mating with the low quality male pays 0.

If the cost for the low quality peacock to grow a big tail (d) is sufficiently large, and if the cost for the high quality peacock to grow a big tail (c) is sufficiently small, then a Nash equilibrium exists in which the high quality male grows a long tail and the low quality male does not and in which the female chooses only to mate with males with long tails.⁶ In this situation, biologists would say that the high quality peacock is signaling his quality to the female.

Under traditional definitions of a signal, if one is to justify such claims about the signaling nature of this game, one must know either the intentions behind the production of the long tail or the evolutionary history of this phenotype. We believe that this distinction can be drawn without reference to either.

Notice that had the peahen been aware—directly—of the peacock's quality, the peacock would not have bothered to grow a long tail. He grows a long tail because it communicates his quality, but only because it does this. In order to make this informal idea somewhat more precise, we introduce a modified version of the peacock game where the peahen is aware of the peacock's type. This game is pictured in Fig. 2.

In this "comparator game," the peahen is aware of the peacock's quality. As a result neither the high nor the low quality peacock will grow a long tail — doing so would expend resources for no purpose. So, although there is a Nash equilibrium in the original game that involves high quality peacocks producing a large tail, there is no such equilibrium in this comparator game where the peahen starts off just as well

⁴ This is a standard story in biology. But it is important to note that it should not be taken as a formal natural history of peacocks and peahens. For simplicity we use teleological language throughout; 'peahens want x ' is our shorthand for 'peahens are selected to ...'

⁵ In biology, the germ of the idea is ascribed to Zahavi (Zahavi, 1975). A game theoretic foundation was later provided by Grafen (1990). In economics, the first models of signaling were developed by Spence (Spence, 1973).

⁶ A Nash equilibrium of a game is a stable state where each individual player is doing the best she can given how the other player(s) are behaving.

³ Skyrms (Skyrms, 2010) is somewhat difficult to categorize as clearly teleosemantic or otherwise. He does not attempt to develop a distinction between signals and cues. He does, however, develop a theory of misinformation and deception which uses properties of the entire population. A signal is roughly speaking misinformation for Skyrms if it normally is used to communicate about state S but in this case is used in a different state.




Fig. 1. Shown in extensive form, a costly signaling game in which a peacock signals his quality to a peahen. Double lines indicate equilibrium play for the signaling equilibrium that exists when $c < 1 < d$. At this equilibrium, the size of the tail conveys reliable information about the quality of the peacock.

informed as the peacock.

This observation identifies what we call an informational move:

First gloss. A move m by player 1 is *purely informational* if it is performed when player 1 knows things that are unknown to player 2, but not performed when player 2 has all of the information that was available to player 1 when player 1 chose m .

We will make this notion mathematically precise in section 5, but first we will turn to another example.

3. Information concealing moves

In their (now classic) discussion of signaling games Cho and Kreps (Cho & Kreps, 1987) describe a scenario known as the beer and quiche game. A traveler finds himself with no choice but to have breakfast at a particularly rough bar. Visitors often end up at this bar for breakfast and the locals occasionally challenge the visitors to a fight. Some of the visitors are surly fellows who would prefer to drink beer for breakfast. Beyond their culinary preferences, surly fellows are dangerous to fight and the locals would prefer to avoid them. Not all visitors are surly, however. Some are wimpy. Wimpy visitors have a preference for quiche for breakfast, and are exactly the sort of people that the locals would like to fight. Importantly like the example of the peacocks before, the locals cannot directly observe whether the visitor is surly or wimpy.

This game is illustrated in Fig. 3. In this game there is a Nash equilibrium where both the wimpy and surly visitors choose to drink

beer. The surly individuals drink beer because they prefer beer. The wimpy individuals, however, cannot afford to choose quiche despite their preference for it, because by doing so they would reveal that they are wimpy and would then be challenged to a fight. Instead, even the wimpy individuals drink beer in order to disguise their fighting abilities.

We can again consider our test for informational moves. Suppose now that the locals can observe directly whether the visitor is surly or wimpy. In this situation (shown in Fig. 4) the surly visitor continues to drink beer – it is his preferred breakfast after all – while the wimpy visitor now consumes quiche. Under our test, proposed in the previous section, the consumption of beer by the wimpy visitor is an informational move while the consumption of beer by the surly visitor is not.

This informational move is somewhat different from the case of the peacocks presented in the preceding section. A peacock's long tail carries something like Grice's non-natural meaning. One might reasonably say that the long tail non-naturally means that the male is of high quality. In the beer and quiche game, the converse occurs. Here the consumption of beer by the wimpy guy is not to convey non-natural meaning, but rather to conceal natural meaning. Drinking beer remains in some respects a communicative act, but of a very different sort from growing a long tail in the peacock scenario. We use the label *information concealing moves* for actions like drinking beer despite preferring quiche.

At this point, it is instructive to compare our notion of “information concealing” moves to other notions of misinformation and deception in




Fig. 2. The peacock "comparator game"




Fig. 3. The beer and quiche game.




Fig. 4. The beer and quiche “comparator game”

games. Skyrms (Skyrms, 2010), for example, defines a signal as misinformation if it raises the probability of a state which is not actual.⁷ Notice that in the beer and quiche game, the action of ordering beer has no effect on the probabilities of the visitor being wimpy or surly. So rather than being misinformation, ordering beer is no-information. This is why we have opted to use the term information concealing rather than deception or misinformation—we believe these are genuinely distinct concepts.⁸

4. Generalizing to non-generic games

The signals we considered in the previous section have production costs that are related to their meanings. High quality peacocks can produce long tails at lesser expense than low-quality birds; surly fellows actually enjoy a breakfast of beer instead of quiche. A short tail cannot come to “mean” that the male is of high quality; eating quiche cannot come to “mean” that a diner is surly.

In their purest form, conventional signals have no relation between meaning and production cost (Guilford & Dawkins, 1995). The signals referred to in the previous sections would probably not be considered (fully) conventional—the structure of the interaction prohibits the reversal of the signals. The paradigmatic case of conventional meaning arises in a Lewis signaling game (Lewis, 1969). As in a costly signaling game, this kind of game has two players, a sender and a receiver. The sender observes the state of the world and sends a signal. The receiver responds to the signal by choosing an act. Unlike the costly signaling

case, there is complete common interest between the sender and the receiver. Moreover, sending a signal is cost-free.

If there are two states, two signals and two acts, the extensive form of the Lewis signaling game is as in Fig. 5, where ϵ_A and ϵ_B are zero. The choices of moves indicated in the figure constitute a signaling system. In this signaling system, the sender chooses signal a in response to state A , and the receiver chooses the appropriate act for that state. Similarly, in state B the sender signals b and the receiver chooses the appropriate response. What makes this kind of signaling purely conventional is that a and b have neither any intrinsic meaning nor any structural properties that prevent their meanings from being reversed. As indicated by the payoffs, the sender might as well choose b in state A and a in state B . As long as the receiver responds correctly, they are as successful as before.

With these payoffs, the Lewis signaling game is a *non-generic* game in extensive form. This means that some of a player's payoffs at terminal nodes are equal. Non-genericity turns out to be a problem for our definition of an informational move. Recall that the basic idea of an informational move is that a player would not make that move in the comparator game. This is strictly speaking not true for the Lewis signaling game. In the comparator game, the receiver is fully informed about the states of nature and, in equilibrium, takes the appropriate act. This makes the sender indifferent between sending one signal or the other. There are, therefore, equilibria in the comparator game where both signals are sent, but this occurs only because of the presence of payoff ties. Should one of the signals possess the slightest of costs, no matter how small, this would no longer be true. As a result, this fact depends critically on something that is unstable to tiny modifications of the game's payoffs.⁹

There is a simple and time-honored solution to this problem. In a first step, it requires us to turn the Lewis signaling game into a generic

⁷ Skyrms is not the only theorist who draws a distinction between informative and misinformative signaling. A similar distinction could be made with other definitions of deception/misinformation.

⁸ Notice that case of outright lying, like telling a friend that his karaoke singing is beautiful, is under our definition a “signal” even if it is a misinformative or deceptive one.

⁹ While some of the earlier games we described are also non-generic, they are so in an innocuous way since this problem does not arise.




Fig. 5. The generic Lewis signaling game.




Fig. 6. The generic Lewis "comparator game"

game. This can be done in the game of Fig. 5 by choosing different values for the ϵ 's. The comparator game of this generic version of the Lewis signaling game is shown in Fig. 6. Now signal b is clearly an informational move by the sender because she would not choose b in the comparator game.

Notice that this idea does not depend on how small the epsilons are that we choose in order to make a generic Lewis signaling game. This suggests the following definition:

Second gloss. In a non-generic game, a move by player 1 is purely informational if there exists a sequence of generic games with payoffs converging to the payoffs of the non-generic game such that player 1's move is informational in each of the generic games of the sequence.

This definition allows us to apply our idea of informational moves to extensive form games even if their payoff structure is not generic. The idea is essentially the same as the one that underlies Reinhard Selten's concept of a *trembling hand perfect equilibrium* (Selten, 1975), where one requires of a Nash equilibrium to be the limit of the Nash equilibria of a sequence of payoff-perturbed games.¹⁰

5. Core theory

We begin by introducing the basics of extensive form games. (Unfamiliar readers should consult any introductory textbook on game theory for more precise definitions.)

A game in extensive form begins with a rooted tree, a special kind of mathematical graph which has a beginning node and has no cycles. A node that is not at the end of the tree is called a decision node and is labeled with a player who moves at that node. "Nature" is treated as a player who makes moves with some fixed probability – these represent non-strategic decisions which influence the outcome of the game. Every node that is at the end of the tree is called a terminal node and is assigned a payoff value for each player – these represent the various ways that the game might end.

Some nodes may be collected together in information sets. These information sets represent ignorance by a player. If nodes n and n' are both in the same information set, we interpret this as indicating that the player is not informed whether she resides at n or n' . There are various constraints on information sets that formally prevent one information set from containing nodes of more than one player, that prevent players from forgetting something they learned earlier in the game, and that prevent players from being ignorant of their own actions.

Because of the complexity of games in extensive form, we will restrict ourselves to a certain class of games known as "action-response games." These games are essentially games where nature chooses from a set of options with a fixed probability. One player is (potentially) given some information about the state chosen by nature. This player can then take a move. A second player (potentially) observes some information about the state and (potentially) some information about the first player's action and then takes an action herself. After this the game terminates. Formally, an action-response game is any game with two players plus a move by nature such that every potential path through the game tree features first a move by nature, second a move by player 1, and finally a move by player 2—no more and no less.

A game G is *generic* if there is a unique payoff at each terminal node (i.e., there are no payoff ties). For the first part of this section, we will restrict ourselves to considering generic games.

We now wish to consider whether a particular move m by player 1 is an informational move in a generic action-response game G . To do so, we must specify how to construct the appropriate comparator game.

Given the formal constraints on information sets alluded to above, the information sets create a partition on each player's decision nodes. Let I_1 be the partition of the first player's decision nodes induced by first player's information sets. Let I_2 be the partition induced on the second player's decision nodes induced by the second player's information sets.

Move m is a move by player 1 in G , and s is the information set where the first player can choose action m . The information set s represents all the information that player 1 has about the move by nature when he has the option to take move m . Let $D(m)$ be the two-element partition of 2's moves such that all moves which come after the information set s in the tree will be in one element of $D(m)$ and the remaining nodes will be in another. Essentially s and $D(m)$ represent the same information; s in terms of what player 1 knows about his own choices and $D(m)$ in terms about what player 1 knows about what player 2 can do. Like s , $D(m)$ represents exactly the information that player 1 has when he has an option to take move m . From the perspective of player 2 the two elements of $D(m)$ represent the state of knowledge: "we are in state s " or "we are not in state s ", where the latter contains no more information than we are in some state other than s . Importantly, $D(m)$ does not represent anything counterfactual about player 1, it does not represent what information player 1 might have had if he had been able to take another action in a different part of the game tree. In other words $D(m)$ represents exactly what player 1 knows is the case at information set s from what he knows is not the case at that information set.

We now construct the comparator game $C(G, m)$ so that the second player knows everything he knew in the original game, plus everything that the first player knows when he makes the move m . To do this, we keep the tree, nodes, and payoffs the same as in G . We also keep the information sets of the first player the same. We alter the information sets of the second player, such that they become the join $I_2 \vee D(m)$, i.e., the coarsest common refinement of I_2 and $D(m)$. That is, we add the knowledge "we are in state s " to player 2's information set in the most minimal way possible (without telling them anything else in any other state).

A consequence of this procedure is that the comparator game is defined relative to a particular move m . When evaluating different moves, different comparator games will be constructed that will have different informational structures. With this in hand, we can now provide a formal definition of an informational move in a generic game:

5.1. Definition

Let G be a generic action response game in extensive form. A move m by player 1 in G is an informational move if m is performed with positive probability in some Nash equilibrium of G but is not performed with positive probability in any subgame perfect Nash equilibrium in $C(G, m)$.

This definition is, again, satisfied by the peacock example. In the original game there is a subgame perfect Nash equilibrium where the peacock grows a long tail, but no subgame perfect Nash equilibrium where the peacock does so in the comparator game. Similarly for the other games we have discussed.

As can be seen from our informal definitions above, this way of understanding informational moves has a certain counterfactual quality to it. One can think of G as representing the situation in the actual world and $C(G, m)$ as representing the closest possible world where player 2 has all the information available to player 1 when player 1 chose m .

With this interpretation in hand, we will now discuss some of the assumptions that underly our definition of informational moves. First, we are presuming that in the actual world we are considering a situation where all players are in a Nash equilibrium when they play this game. Undoubtedly there are many situations in this world where players are out of equilibrium. Because of the relatively weak

¹⁰ We can create one sequence that makes signal A informational and another sequence that can make signal B informational. So, on our definition both come out as informational in the original non-generic game.

assumptions of Nash equilibrium, in order to be out of equilibrium the players must either be irrational or have incorrect assumptions about one another's behavior. In the biological context, to be out of equilibrium means that the population must be subject to change due to natural selection.

We do not think the notion of informational moves in out-of-equilibrium situations are particularly helpful, because it might be the case that player 1 *thought* she was sending information to player 2, but player 2 thought otherwise. Here we believe there is not a clear fact of the matter regarding the informational status of the situation.

Second, we argue that the game $C(G, m)$ characterizes the strategic situation facing the players in the closest possible worlds where player 2 has all the information available to player 1 when she takes move m . Here we equate “closest possible world” with minimum modification of the game tree. While there might be situations where this isn't true, we believe that it should be uncontroversial for most situations of interest.

Finally, we presume that in the closest possible world $C(G, m)$ players are playing a subgame perfect Nash equilibrium. Subgame perfect equilibria are a refinement of Nash equilibria where players are not making non-credible threats—they are not promising to take moves that, if forced to, would be irrational.

We utilize this more restrictive solution concept because in a number of cases that we consider the informational move might continue to be played but only because player 2 is threatening to harm both herself and the other player if the move is not taken. We do not believe that such irrational commitments are plausible, nor are they helpful in uncovering the phenomena with which we are concerned. That we do not use this more restrictive notion when considering play in the actual world should not be taken as a negative comment against this equilibrium notion, but rather a preference for mathematical generality.

Prior to turning to our distinction between information concealing and revealing moves, we must provide one more definition to handle non-generic cases, like the Lewis signaling game.

5.2. Definition

Let G be a non-generic action-response game in extensive form. A move m in G is said to be an informational move if there exists a sequence of generic games G_n such that $G_n \rightarrow G$ and m is an informational move in every G_i .

As in the definition of the comparator game, the sequence of games may be different for different moves m and m' .

6. Information revealing or information concealing?

In our discussions of the peacock game and of the beer and quiche game we found that some informational moves, such as the peacock's big tail, can be information revealing while other informational moves, such as the wimp's beer breakfast, can be information concealing. In this section we provide a formal mechanism for determining whether a given move is information revealing, information concealing, or neither.

The basic intuition is as follows. A move is information revealing if it gets you to someplace you wouldn't go in the comparator game; a move is information concealing if it avoids going someplace you would go in the comparator game.

To formalize this, consider an action response game G , a move by player 1, m , in G , and the comparator game $C(G, m)$. Let h be the natural map from nodes in $C(G, m)$ to nodes in G . Let f be a map from information sets in $C(G, m)$ to information sets in G based on h . Suppose node n is in information set S of the comparator $C(G, m)$. Then $h(n)$ is a node of game G in some information set T in G , and we set $f(S) = T$. (Since $C(G, m)$ has exactly the same information sets as G except for the moves following m , such a function f is always well-defined.)

Let m be an informational move. Let $E_G(m)$ be the set of all information sets of G reached in a Nash equilibrium of G where m is played. Let $E_C(m)$ be the set of all information sets in $C(G, m)$ reached in some subgame perfect Nash equilibrium of $C(G, m)$.

6.1. Definition

- A move m is *information revealing* if and only if (1) m is informative and (2) $f(E_C(m)) \subset E_G(m)$.

Consider the example of the peafowl with a focal move, m , where the peacock grows a long tail. In the peafowl game there are two information sets, the set where the peacock grows a long tail the set where he does not. There is a Nash equilibrium, the signaling equilibrium, where both information sets are reached. As a result, $E_G(m)$ equals both information sets.

However, in the comparator game $C(G, m)$ the nodes for the peahen that occur at the top of the picture, where the peacock grows a long tail, will never be reached in any subgame perfect Nash equilibrium of $C(G, m)$. As a result $E_C(m)$ will not contain those nodes and the information set at the top of the picture will not be a member of $f(E_C(m))$.

In order to be an information revealing move, we require that nodes be reached in the original game that would not be reached in the comparator. This captures the idea that information is being revealed rather than concealed. We now turn to that latter type of informational move.

6.2. Definition

- A move m is *information concealing* if and only if (1) m is informative and (2) $f(E_C(m)) \supset E_G(m)$.

To illustrate this definition, consider the beer and quiche game. In the original game, G , the only information set that is reached is the one where the visitor orders beer for breakfast. But, in the comparator game, both beer and quiche are ordered. As a result, ordering beer is regarded as an informational concealing move in the beer and quiche game.

7. Conclusion

Biologists, economists, and researchers in a range of other disciplines commonly study signaling games, using the machinery of game theory to model communication. But what makes a signaling game a signaling game? Or in other words, where in a signaling game is the signal? Any biologist would answer this by appealing to whatever story he or she told to motivate the analysis – “the peafowl game is a signaling game because a high quality peacock has an incentive to communicate his quality to a potential mate.” But this is not a game-theoretic answer. Game theory is not about the stories that one tells alongside one's model; it is the about the analysis of formally defined mathematical objects called games. These objects are fully described by the extended form (or even the normal form) of the game. This poses a problem: it is only meaningful to talk about signaling games if some games are signaling games and others are not. And if this is the case, we obviously have to be able to determine which is which from the extended form game itself, not from the story that someone spins to go along with it.

This paper presents the start of a theory of informational moves which avoids the use of intentions, teleology, or motivational stories. While not yet fully general, our hope is that this approach might evolve into a fully general theory of signaling in social interaction. This would enable a more clear articulation of the distinction between signals and other types of biological and economic phenomena, and may provide the groundwork for an alternative theory of meaning. Much work in

this direction remains, but we hope that this provides an important starting point.

CRediT authorship contribution statement

Carl T. Bergstrom: Conceptualization, Writing - original draft, Writing - review & editing. **Simon M. Huttegger:** Conceptualization, Writing - original draft, Writing - review & editing. **Kevin J.S. Zollman:** Conceptualization, Writing - original draft, Writing - review & editing.

Acknowledgements

This work was supported by National Science Foundation grant EF-1038456

References

- Birch, J. (2014). Propositional content in signalling system. *Philosophical Studies*, 171(3), 301–318.
- Cho, I.-K., & Kreps, D. M. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2), 179–222.
- Dretske, F. (1981). *Knowledge and the Flow of information*. Cambridge: MIT Press.
- Gneezy, A., Gneezy, U., Riener, G., & Nelson, L. D. (2012). Pay-what-you-want, identity, and self-signaling in markets. *Proceedings of the National Academy of Sciences of the United States of America*, 109(19), 7236–7240.
- Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144, 517–546.
- Grice, H. P. (1957). Meaning. *Philosophical Review*, 66(3), 377–388.
- Guilford, T., & Dawkins, M. S. (1995). What are conventional signals? *Animal Behaviour*, 49, 1689–1695.
- Harms, W. F. (2004). *Information and meaning in evolutionary processes*. Cambridge University Press.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge: Harvard University Press.
- Maynard Smith, J., & Harper, D. (2003). *Animal signals*. Oxford: Oxford University Press.
- Millikan, R. G. (2005). *Language: A biological model*. New York: Oxford University Press.
- Peirce, C. S. (1867). On a new list of categories. *Proceedings of the American Academy of Arts and Sciences*, 7.
- Scott-Phillips, T. (2008). Defining biological communication. *Journal of Evolutionary Biology*, 21(2), 387–395. <https://doi.org/10.1111/j.1420-9101.2007.01497.x>.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1), 25–55.
- Shea, N., Godfrey-Smith, P., & Cao, R. (2016). Content in simple signalling systems. *The British Journal for the Philosophy of Science*, 1–27 0(March).
- Simons, M., & Zollman, K. J. S. (2019). Natural conventions and indirect speech acts. *Philosophers' Imprint*, 19(9).
- Skyrms, B. (2010). *Signals: Evolution, learning and information*. New York: Oxford University Press.
- Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.
- Stegmann, U. E. (2009). A consumer-based teleosemantics for animal signals. *Philosophy of Science*, 76(5), 864–875.
- Zahavi, A. (1975). Mate selection – a selection for a handicap. *Journal of Theoretical Biology*, 53, 205–214.