Sender-Receiver Systems Within and Between Organisms

Peter Godfrey-Smith

City University of New York

Drawing on models of communication due to Lewis and Skyrms, I contrast sender-receiver systems as they appear within and between organisms, and as they function in the bridging of *space* and *time*. Within the organism, memory can be seen as the sending of messages over time, communication between stages as opposed to spatial parts. Psychological and genetic memory are compared with respect to their relations to a sender-receiver model. Some puzzles about "genetic information" can be resolved by seeing the genome as a cell-level memory with no sender.

1. Introduction

The investigation of meaning and representation is being transformed by new work that studies signaling games of the kind originally described by David Lewis (1969), but treats them from an evolutionary point of view.[1] This approach was introduced by Brian Skyrms (1995, 2010). Lewis's model was designed to apply to communication between two human agents. But as Skyrms and others have shown, the model has more general application. It can be applied to communication between nonhuman organisms, from mammals to bacteria, and it can be used to address questions about communication between parts of an organism as well. This paper looks closely at the application of "sender-receiver" models of this kind to within-organism interactions.[2]

I first discuss the structure of the models, looking especially at the forms of control and coordination that signaling achieves. I then introduce a distinction between clear cases of sender-receiver systems and partial or *marginal* ones – systems with only an approximate fit to the sender-receiver model. The framework is then applied to

---

[1] This paper was presented at PSA 2012, San Diego, as part of the symposium "Signaling Within the Organism," with Nicholas Shea, Rosa Cao, Brett Calcott, and Rory Smead.

[2] Harms (2004) is an early investigation of the significance of Skyrms's work on the Lewis model for the philosophy of mind. The present discussion is also indebted to earlier work applying information-theoretic and evolutionary concepts to these areas (Dretske 1981, Millikan 1984).

memory, understood as a general functional category: memory is communication over time. I then move to a comparative treatment of cases. The main examples are neural and genetic memory, two ways organisms retain useful information over time. In the genetic case, where there is consensus about basic mechanisms, theoretical debate concerns about how these mechanisms relate to notions of information and representation. In the neural case there is more uncertainty about mechanisms; I discuss how the model relates to neural memory on both orthodox and some unorthodox conceptions. In the genetic case I propose a partial resolution of recent debates.

2. The Sender-Receiver Approach

The aim of Lewis's 1969 treatment was to understand "conventional" communication, especially in response to Quine's skepticism about convention. In the Lewis model we assume two rational agents. One is a "communicator" or (as I will say) *sender* who can see the state of the world but cannot act except to produce signals, messages, or signs of some kind. (I will use the term "sign" broadly here.) The "audience" or *receiver* can only see these signs, but can act in a way that affects them both. Lewis assumed *common interest* between the two agents: they have the same preferences for what they want done in each state of the world. Then if the sender sends distinctive signs in each state of the world, and the receiver uses these signs to prompt appropriate actions, such a "signaling system" can be a *Nash equilibrium* – neither side has any incentive to change (unilaterally) what they are doing.

More formally, the sender and receiver each implement rules: f_S for the sender, f_R for the receiver. The sender's rule maps states of the world to signs; the receiver's rule maps signs to acts. Composed, the result is a function from states to acts, F, which may or may not correspond to the preferred mapping for both agents, F^*. When the preferred mapping between states and acts is implemented by the combination of sender's and receiver's rule, that state is a Nash equilibrium.

A simple case like this, with tight coordination between sender and receiver, can be described as achieving a division of labor: the sender acts as the system's eyes, the receiver as hands. This configuration can also be seen as a solution to a particular problem of control: pairing appropriate behavior with states of the world, dealing with the environmental contingency of effective action.

Lewis assumed that the sender and receiver are rational agents linked by common knowledge. Skyrms (1996, 2010) gave an evolutionary, naturalistic treatment of the model: rational choice was replaced by natural selection and, in further work, by simple forms of learning. These are all "selection" processes in a broad sense of the term. Many differences of detail between these processes have arisen with respect to how they bear on sender and receiver behaviors (Skyrms 2010), but the overall outcome has been that signaling of roughly the kind modeled by Lewis can arise and be maintained by many different selection processes.[3]

The Lewis model also has connections to information theory. When Claude Shannon (1948) gave a quantitative treatment of the transmission of information, he set out by imagining a *transmitter* and *receiver*, connected by a *channel* over which signals pass. He introduced measures of the amount of uncertainty associated with a variable – Shannon's *entropy* – and various measures of the degree of association between two variables – the extent to which the value of one variable predicts the value of the other. One of these measures is now called the *mutual information*.[4] Though Shannon's model was set up by imagining something like a sender and a receiver, those roles are not essential to his framework, and his core theoretical concepts can be applied in many other contexts. Mutual information between variables is ubiquitous in physical systems; the Lewis model describes one way that it can come to exist: by the shaping of senders' and receivers' behaviors. The sender is directly responsible for the informational relations between signs and states of the world – responsible for whether the signals sent bear any systematic relation to the states perceived – but the sender's payoffs, and hence their choices, are dependent on what the receiver will do with those signs. Sender and receiver

---

[3] For an evolutionary treatment of the Lewis model see also Huttegger, Skyrms, Smead and Zollman (2010).

[4] The *Shannon entropy* of a discrete random variable X which has possible values x_1, x_2, x_i... is $H(X) = -\sum_i P(x_i) Log_2 P(x_i)$. The *mutual information* between two variables X and Y is $H(X)+H(Y)-H(X,Y)$. Here $H(X,Y)$ is the "joint entropy" of X and Y, the entropy of the distribution of combinations of X and Y values.

behaviors coevolve, and signs have their informational properties shaped as a consequence.

When the sender's rule maps states to signs one-to-one, for example, the signs carry as much information about the world as is possible. But the sender will only perform such a mapping if the receiver will do something useful – from the sender's point of view – with the information supplied. If sender and receiver disagree entirely about which acts they want performed in each state of the world, then any information about the world contained in signals is likely to be exploited by the receiver. The sender, it seems, should then stop sending informative signals and the receiver, in turn, should stop paying attention to anything the sender says. Recent computational work has uncovered cases where informative signaling can exist at equilibrium in situations where sender and receiver disagree entirely in their preference ordering over actions in every state (Godfrey-Smith and Martinez forthcoming). These cases are rare, and the same investigation finds that as interests diverge in a Lewis model, informative signaling generally becomes more precarious. The result of this divergence may be a refusal to send informative signs, cycles of attempted exploitation, or chaotic dynamics (Wagner 2012).

Information in a sign-like object can also be useable by a receiver or interpreter without a sender being involved at all. Biologists use the term "cue" for unsent or naturally occurring signs that provide useful information. Clouds are cues of the upcoming weather. From the point of view of the sender-receiver model, these are cases where there is a mechanism of "production," in a broad sense, behind a sign, but that mechanism is unaffected the consequences of the sign's use by a receiver. The factors responsible for the physical association between clouds and rain do not care how clouds are interpreted; the feedback that is central to the sender-receiver model is absent.

A smart agent will make use of both sent and unsent signs. These will have differences in role, however. Here is a plausible principle: in cases where there is no sender producing signs in a way affected by the receiver's actions with some degree of common interest, there is no way for sign-like objects to become complicated and elaborate. Complexity in signs cannot *accumulate* by a selection process. This fact does not limit the usefulness of cues, provided that the observer is smart enough. A smart observer can infer a vast amount from traces left fortuitously. (Consider Laplace's demon.) But such use of cues, it seems, requires powerful inferential capacities; it cannot be done by a simple, mechanical reader device. This principle will revisited, however, later in this paper.

This consideration of unsent signs or cues illustrates a more general feature of the application of the Lewis model to empirical systems. The Lewis model, and later developments of it, are highly idealized. The match between the model and any empirical system will be approximate. Some systems have a rather close match to the model – consider bee dances – while others have a more tenuous or partial one. Below I make use of a distinction between clear or paradigm cases of a sender-receiver system and partial or marginal ones.[5] Many empirical features are relevant to this distinction, but a particularly important one is whether the empirical system contains distinct objects that have a reasonably good fit to the roles of sender and receiver. The entities playing these roles must be distinct from the signs lying between them, must be able to implement rules of the kind relevant to the model (discussed above), and must also be susceptible to adaptive change as a consequence of the pairings of actions with states. "Cue"-like cases, discussed above, illustrate one violation of the model's assumptions – the sender or producer is not adapting in response to the receiver's actions. But the goodness of fit of a real-world mechanism to the roles specified by the model is a matter of shades of grey. The sender, for example, might be adapting on a slower scale than the receiver, or there might be a partial "entanglement" of entities the model treats as separate (e.g., sign and receiver).

A further issue concerning the model's application also arises. As discussed above, the relations between a Lewisian sender and receiver enable the solution of a particular problem of control. The problem is the need to pair actions with states of the world, and a solution is achieved by the sender tracking those states and providing guidance to the receiver. Looking at a sender-receiver system of this kind as a whole, the role played by signaling is *act-to-state coordination*.

---

[5] This applies a methodology used also in Godfrey-Smith (2009).

Not all cases of communication, even simple ones, fit this category.[6] Consider a "cox" in a boat calling the stroke for two rowers. The sender has no private knowledge of the state of the world, and it does not matter when exactly the two rowers row, as long as they row at the same time. (The cox is like a "central pattern generator" in the brain.) Here, the point of communication is to coordinate the actions of the rowers. This can be called *act-to-act* coordination, in contrast with *act-to-state*. The difference concerns whether or not there is an exogenous state, not controlled by an agent within the communicative set-up, with which actions need to be coordinated. The "cox" case is one where (in the simple version described) there is no such state. Another example, featuring partial common interest, is the "battle of the sexes" game (Farrell and Rabin 1996). Many cases combine elements of both kinds of coordination, and some might have an uncertain status. Though the Lewis model in its original form concerns the coordination of acts with exogenous states, it can be applied to some cases that concern the coordination of acts with acts; if instead of a cox, one of the *rowers* calls the stroke, then he or she may have private knowledge of the timing of their upcoming stroke, and the Lewis model can be applied, in at least a preliminary way, by treating the sender's intention to row as a "state." Even here, though, the control of the "state" by an agent raises further possibilities not covered in the Lewis model.

In the remainder of this paper, I use the phrases "sender-receiver system" and "sender-receiver model" both for cases that follow the Lewis format and some that have a different structure, but where the crucial relationships in the system still concern the coevolution of sender and receiver behaviors. "Sender-receiver *system*" refers to an empirical systems as opposed to their models. "Senders and receivers" include writers and readers, producers and consumers (Millikan 1984), inscribers and interpreters, publishers and subscribers, of many kinds. As Lewis said in 1969, if we know what the sending and interpreting rules in a system are and why they are the way they are, nothing has been left out of an understanding of the semantic properties of signs in that system. The way to understand sign-related phenomena is to investigate the reciprocal shaping of sender and receiver behaviors, the two-sided phenomenon of sign use.

---

[6] In developing the material in this paragraph I am indebted to discussion with Kevin Zollman.

3. Within the Organism

Sender-receiver systems arise in a wide range of contexts. At least three kinds of generality can be distinguished. First, the model can be applied to signaling *between* and *within* organisms. Second, a range of selection processes operating on different time-scales can shape sender and receiver behaviors, including evolution by natural selection, reinforcement learning, differential imitation, and deliberate choice. Third, signs can be used by senders and receivers to bridge both *space* and *time*. This gives the model a connection to *memory*, both in psychological and other forms.

The most conspicuous communication-like activity within organisms is interaction between neurons, a case noted by Skyrms (2010). Hormonal systems are another example. The entire sensorimotor axis of an organism can also be seen as realizing a sender-receiver structure, in at least a rough way; the problem of coordinating acts with the world is solved by coordinating receptive and active parts of the organism (recall a metaphor used in the previous section; a Lewisian sender acts as eyes, and receiver as hands, in a cooperating system).[7]

These are cases of specialization of the parts of an organism, where "parts" are spatial parts. But as well as coordination of spatial parts, there is the coordination over time. Stages of an object (sometimes seen in metaphysics as "temporal parts") can also function as sender and receiver. Present stages often have access to information that will be useful to the guidance of action by future stages, and present stages can lay down marks to be consulted later; memory is sending messages over time.

Those points enable an initial application of the sender-receiver model to within-organism relations, but as noted above, the set-up characterized by the sender-receiver model appears in both clear and marginal cases. In different settings, different kinds of partial cases arise, as a consequence of local raw materials and the evolutionary processes operating. In the between-organism context, complications and departures from a Lewisian structure arise especially from the divergent interests and multifarious agendas of communicating agents. In the within-organism context, common interest is not as

---

[7] See also O'Connor (forthcoming) for the within-organism application of the sender-receiver model focused on the sensorimotor axis and the evolution of categorization.

much of an issue, though there are exceptions (Haig 2008). Complications arise especially from the entanglement of roles treated as distinct in the model.

In the simplest neural systems (Jékely 2011), individual cells are plausible senders and receivers and neurotransmitters are signals. But in more complex neural systems – especially as interneurons proliferate – the roles become entangled (Cao 2012). Neurons collectively mediate between senses and effectors, but each neuron is upstream of some and downstream of others, and indirect recurrent connections also arise. A whole neuron mediating between two others is not passive in the way a Lewisian sign is; its activity is not a mere means for one neuron to affect another. In the within-organism context, where conflict of interest is limited, the active capacities of neurons enable them to take on more elaborate roles in control systems.

The same considerations apply in the temporal dimension. One way for memory to operate, a way that makes it a clear case of a sender-receiver system, is to have a *write-read* structure. This kind of memory is familiar from external mnemonic devices such as notes and diaries. It is also is seen in simpler form in a Turing machine, where marks are made on, and then read from, a tape. The same applies to memory in ordinary modern computers. According to Danziger (2008), an inscription model has been, from the time of Plato, the "master metaphor" in the Western tradition of theorizing about psychological memory. Many contemporary neuroscientists, however, make a point of claiming that a write-read structure is *not* in general how memory works in our brains. Memory is primarily achieved, according to mainstream views, by modifying the weights of synapses (modifying the way the firing of one neuron excites or inhibits another). Once a synaptic weight has been changed, this affects the later operation of neural networks that contain the synapse, but this change is not one in which a reader device gains access to a new sign; it is as much a change to the processing machinery as a change to what they process. This entanglement is expressed explicitly by Christof Koch, for example: "memory is everywhere, intermixed with the computational elements" (1999, 471). Many contemporary neuroscientists would probably agree, and would see the message of recent work as telling against a write-read view of memory.

Against this background, Randy Gallistel and Adam King (2010) have recently argued that the move away from a write-read model for neural memory has been a mistake, one based on spurious arguments about biological plausibility and optimistic extrapolation from simple phenomena of synaptic plasticity. Gallistel and King argue that given what brains (including simple non-human ones) can evidently do, they must contain an "addressable read/write memory," one that is more like computer memory than anything envisaged in current neuroscience. Synaptic weights are a possible, but unlikely, location for this form of memory, and they suspect that an entirely different form of information storage in the brain has yet to be recognized.

John Donohoe, a defender of the mainstream view, wrote a review of Gallistel and King's work in which he said that what we have learned is that the brain contains a "write-only" memory system (2010). Gallistel replied that if the brain wrote something but could not read it, that would surely be pointless. Donohoe's view might be better expressed by saying that, according to mainstream views, the brain contains a *write-activate* memory. Evolved neural machinery has the function of introducing marks or traces into the brain as a result of experience, but these marks have useful effects on behavior without being read. From the point of view of the sender-receiver model, this is a debate about whether neural memory is a clear realization of a sender-receiver structure or a more marginal one, because of the entanglement of the roles of sign and receiver.

4. Genetic Information

I now turn to another case. Informational and communication-related concepts have become important parts of modern biology, especially in the description of gene action and the evolutionary role of genes. The status of these descriptions has been much debated. Especially ambitious claims have been advanced by G.C. Williams (1992). Williams holds that in the light of modern biology, we should recognize a "codical domain," comprised of information, in addition to the familiar domain of material objects. Evolutionary change takes place in both domains. Above I discussed a view of neural memory advanced by Gallistel and King. They defend their *un*orthodox view of neural memory by appealing to the importance of write-read mechanisms in the *orthodox* view of genetic systems; the genetic case shows that evolutionary processes can construct sophisticated write-read systems at the subcellular level, so there can be no argument based on biological plausibility against this model when applied to the brain.

On the other side of the debate about genetics is a tradition of argument that sees enthusiasm for the use of informational concepts in biology as indicating a slide into illegitimate preformationist patterns of explanation, and perhaps a subtle form of genetic determinism (Oyama 1985, Francis 2003; arguments surveyed in Godfrey-Smith and Sterelny 2007). Intermediate positions have also been advanced. Godfrey-Smith (2000) and Griffiths (2001) argued that a narrow sense of "genetic coding" has a genuine theoretical role in biology: DNA sequences code for the primary structure of protein molecules. This defensible use of a semantic concept in cell biology is surrounded by many vaguer analogies and over-extensions. A different kind of intermediate position has been developed by Arnon Levy (2011), who argues for a strictly "fictionalist" view of informational and semantic concepts in cell and developmental biology, but one that emphasizes the positive heuristic role these concepts play in organizing knowledge of biological mechanisms.

Some valuable recent papers have considered this issue explicitly within a sender-receiver model, or a close relative of it: Shea (2007, 2012), Bergstrom and Rosvall (2009). The aim of this work is to justify quite strong versions of the idea that gene sequences represent phenotypes (Shea), or carry information in rich sense (Bergstrom and Rosvall).[8] Here I present a different view, applying the version of the sender-receiver model outlined above. I then make contrasts with the Shea-Bergstrom-Rosvall approach. I argue that genetic mechanisms constitute a distinctive kind of marginal case of a sender-receiver system, one where a particular aspect of the paradigmatic structure is absent or attenuated.

The best starting point for a view of gene action that takes concepts of coding and representation seriously, as acknowledged in much of the earlier literature, is the reality of the "reading" step. Genes have a rather clear reader mechanism, the transcriptional and translational machinery. The next move is not so clear. My approach is to make an explicit link to the idea of memory. This link was made very early in the discussion, by David Nanney. He noted in 1958 that there are two distinct tasks that cells manage to perform:

---

[8] I have discussed the details of these proposals elsewhere. For Shea, see Godfrey-Smith and Sterelny (2007), for Bergstrom and Rosvall, see Godfrey-Smith (2011).

> On the one hand, the maintenance of a "library of specificities," both expressed and unexpressed, is accomplished by a template replicating mechanism. On the other hand, auxiliary mechanisms with different principles of operation are involved in determining which specificities are to be expressed in any particular cell. … [These] will be referred to as "genetic systems" and "epigenetic systems." (1958, 712)

This is a helpful way of looking at the situation, now as well as then. A genome can be seen as containing a memory of the structure of useful protein molecules. This is not the only thing a genome does, as a genome also participates in the regulation of its own expression (Nanney's second task). But the memory role is one genuine role. Specifically though, a genome is a *cell*-level memory, not a memory for a whole multicellular organism, or a population. The reading process that the genome is involved with is an intra-cellular process, and the result of the reading is production of a protein molecule (or an RNA molecule with another role in the cell). Cells read their genes, and cells persist and divide, carrying their internal library with them. Organisms are the results, the upshots, of what the cells are doing.

So far, this seems to fit well with a temporally organized version of a sender-receiver structure. But this impression comes from looking specifically at the reader or receiver side. Who then is the sender? Which sender's activities make it the case that the cells contain DNA sequences worth reading? The short answer is that there is no sender.

The longer answer is given by revisiting points made in section 2. I noted there that reading or interpretation, in a broad sense, is not confined to messages that have a sender. Unsent signs or "cues," in the biologist's sense, can be useable. But in general, I suggested, without a sender shaped by selection there is no way for a complicated and elaborate message worth reading to arise. More exactly, the relationship is one between sign complexity and reader complexity. All natural events leave traces that a sophisticated reader or interpreter can make use of by means of inference, and this can include natural traces that are very complicated. (We "read" the chemical composition of layers in the earth's crust, in this sense.) If the reader is not sophisticated, but is simple and inflexible, then in general the only unsent "signs" that can usefully be read are simple ones. An example is the use of day-length by plants as a cue of the change in seasons.

Turning again to genes, we find that the reader mechanism present here is a simple and inflexible one, but the genome being read is a complex object whose interactions with the reader mechanism are very fine-grained. It would then be natural for Shea, Bergstrom, and Rosvall to argue that the intricate design and the amount of useful information contained in DNA shows that it was *made for* an information-carrying role, and genetic messages do have a sender. Perhaps the sender is the previous generation.

However, the embedding of DNA within evolutionary processes gives rise to a special case here. The genetic sequences present at any time are shaped by evolution by natural selection. Though evolution is one process that can *build* a sender mechanism, evolution *itself* is not a sender mechanism; it is just something that happens. Mutations happen; differential survival and reproduction happen. As a result of how they happen, DNA sequences are made into libraries of useful specificities (in Nanney's phrase), and also into components in control systems that regulate how the library is accessed.

The replication of DNA is part of the story here, and in many organisms, so is recombination. How do they fit in? Replication itself is more closely related to persistence; it is not analogous to inscription in the sense of the sender-receiver model. In DNA replication, a sequence is preserved, with slight modifications. DNA replication is not the creation of a sequence from some other source. Recombination has a special role, because it does give rise to sequences which did not exist before. But there is no writing step in the sense of the model; there is no mechanism which executes a mapping from states to signs (f_S, from section 2), or a similar process whereby a DNA sequence is produced from something else by the operation of an evolved mechanism. Genetic material is shaped, but this occurs by means of mutation and selection, a process in which some sequences are retained and others are lost, and those that are retained are subject to mutation and recombination. The *retention* of genetic information, especially across cell division, is an adapted process, as is its use by a reader; its *inscription* is not.

In short, DNA's evolutionary embedding makes it an unsent sign, something worth reading without having being written, even though it is very complex. Rather than a *write-read* memory system, genetic memory systems comprise an *evolve-read* system. If so, the contrasts between genetic memory and neural memory (as conceived on the mainstream view) can be represented as in Figure 1 .

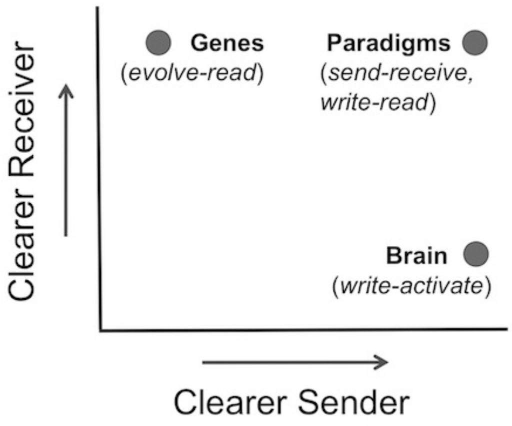

Figure 1: Contrasts between a paradigm sender-receiver system and two empirical cases involving memory. The axes represent the extent to which the system includes a clearly demarcated sender (/receiver) that fits the specifications of a sender-receiver model. The point for neural memory systems represents how this case is understood on a mainstream view, which is contrasted with an unorthodox (upper right) view in section 3.

If this view is right, DNA does not contain a message, a content, of the kind seen in signs with co-adapted senders and receivers. But it is still possible to ask what is being read *from* it, as we could with a cue. What sort of involvement of the DNA sequence with something external to it makes the "message" worth reading? In answering this question we should be guided by facts about the problem of control being handled. The main problem being solved is retention of the capacity to make complex polypeptide molecules. This is analogous to remembering *how*, remembering what to do, not remembering facts (see also Stegmann 2005). But whereas some gene products are necessary or useful to organisms in a wide range of circumstances, some have utility that is specific to some environments rather than others. In the framework of section 2, reading the genetic message enables a mixture of act-to-act and act-to-state coordination.

A brief comparison can also be made to epigenetic systems in cells, systems more often referred to explicitly as "memory." Here there is, in at least some cases, a clear write-read structure. Marks are made by specialized mechanisms on chromosomes (cytosine methylation is a good example here), and these marks then affect transcription. Further discussion of this case will make the comparative treatment richer.

Stepping back from these details, sender-receiver systems have two roles in a complete explanation of semantic phenomena. First, sender-receiver systems are a natural kind, something that nature builds repeatedly in different contexts and at different scales. Specifically, it builds both clear cases and various kinds of marginal or attenuated cases. The second role these systems have concerns our habits of interpretation as observers. Sender-receiver interactions are so familiar in everyday life that a rich set of habits of description have arisen around them; we are used to *talking* about sender-receiver systems as well as living within them. The role of interpreting observer here is distinct from the role of a receiver within a sender-receiver system being observed; the observer is looking on from outside. As a result of our social embedding, we tend to describe all sender-receiver systems, even marginal ones, in ways derived from our habits of dealing with the rich symbol systems of interpersonal communication.

Both neural and genetic systems contain structures with at least a partial match to a sender-receiver model. In the neural case, there is ongoing empirical debate about how good the match is. In the genetic case, there is more consensus about mechanisms themselves. In cases of partial match to the model like these, observers – scientists and philosophers – often find themselves applying habits of interpretation to a marginal case that are only suited to a clearer case. The "strained" character of the resulting descriptions gives rise to wholesale fictionalist attitudes like Levy's (2011). The view I have argued for here utilizes concepts of approximation rather than fiction, though to describe a marginal case in a way that is only literally applicable to a clear case could be seen as fictionalizing. One part of the overall story is an account of how nature gives rise to different variants on the sender-receiver structure in different contexts; another is understanding how these variants interact with the evolving interpretive habits of observers.

References

Bergstrom, Carl, and Martin Rosvall. 2010. "The Transmission Sense of Information." *Biology and Philosophy* 26:159–176.

Cao, Rosa. 2012. "A Teleosemantic Approach to Information in the Brain." *Biology and Philosophy* 27:49–71.

Danziger, Kurt. 2008. *Marking the Mind: A History of Memory*. Cambridge: Cambridge University Press.

Donohoe, John. 2010. "Man as Machine: A Review of *Memory and the Computational Brain*, by C.R. Gallistel and A.P. King." *Behavior and Philosophy* 38:83-101.

Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

Farrell, Joseph, and Matthew Rabin. 1996. "Cheap Talk." *Journal of Economic Perspectives* 10:103-118.

Francis, Richard. 2003. *Why Men Won't Ask for Directions: The Seductions of Sociobiology*. Princeton: Princeton University Press.

Gallistel, Charles R., and Adam King. 2009. *Memory and the Computational Brain: Why Cognitive Science will Transform Neuroscience*. New York: Wiley-Blackwell.

Gallistel, Charles R. 2010. "Response to Donahoe Review." *Behavior and Philosophy* 38:103-111.

Godfrey-Smith, Peter. 2000. "On the Theoretical Role of 'Genetic Coding.'" *Philosophy of Science* 67:26-44

Godfrey-Smith, Peter. 2009. *Darwinian Populations and Natural Selection*. (Oxford: Oxford University Press)

Godfrey-Smith, Peter. (2011). "Senders, Receivers, and Genetic Information: Comments on Bergstrom and Rosvall," *Biology and Philosophy* 26: 171-181.

Godfrey-Smith, Peter and Kim Sterelny .2007. "Biological Information," *The Stanford Encyclopedia of Philosophy (Winter 2007 Edition)*, E. N. Zalta (ed.), <http://plato.stanford.edu/entries/information-biological/>.

Godfrey-Smith, Peter, and Manolo Martinez. forthcoming. "Communication and Common Interest." Under review.

Griffiths, Paul. 2001. "Genetic Information: A Metaphor in Search of a Theory." *Philosophy of Science* 68:394–412.

Haig, David. 2008. "Conflicting Messages: Genomic Imprinting and Internal Communication." In Peter d'Ettorre and David Hughes (eds.), *Sociobiology of Communication: An Interdisciplinary Perspective.* Oxford: Oxford University Press.

Harms, William. 2004. "Primitive Content, Translation, and the Emergence of Meaning in Animal Communication." In D. K. Oller & U. Griebel (Eds.), *Evolution of Communication Systems: A Comparative Approach*. Cambridge: MIT Press.

Huttegger, Simon, Brian Skyrms, Rory Smead and Kevin Zollman. 2010. "Evolutionary Dynamics of Lewis Signaling Games." *Synthese* 172:177-191.

Jékely, Gaspar. 2011. "Origin and Early Evolution of Neural Circuits for the Control of Ciliary Locomotion." *Proceedings of the Royal Society B*: 278:914–922

Koch, Cristoph. 1999. *Biophysics of Computation: Information Processing in Single Neurons*. New York: Oxford University Press.

Levy, Arnon. 2011. "Information in Biology: A Fictionalist Account." *Nous* 45:640–657.

Lewis, David K. 1969. *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.

Millikan, Ruth G. 1984. *Language, Thought and Other Biological Categories*. Cambridge MA: MIT Press.

Nanney, David. 1958. "Epigenetic Control Systems." *Proceedings of the National Academy of Science*s 44:712-717.

O'Connor, C. forthcoming. "Evolving Perceptual Categories."

Oyama, Susan. 1985. *The Ontogeny of Information*. Cambridge: Cambridge University Press.

Shannon, Claude. 1948. "A Mathematical Theory of Communication." *The Bell System Technical Journal* 27:379–423, 623–656.

Shea, Nicholas. 2007. "Representation in the Genome, and in Other Inheritance Systems." *Biology and Philosophy* 22:313–331.

Shea, Nicholas. 2012. "Two Modes of Transgenerational Information Transmission." In *Signaling, Commitment, and Emotion* (ed. Richard Joyce, Kim Sterelny, and Brett Calcott). Cambridge MA: MIT Press.

Skyrms, Brian. 1996. *Evolution of the Social Contract*. Cambridge, MA: Cambridge University Press.

Skyrms, Brian. 2010. *Signals: Evolution, Learning, and Information*. New York: Oxford University Press.

Stegmann, Ulrich. 2005. "Genetic Information as Instructional Content." *Philosophy of Science* 72:425-443

Wagner, Elliott. 2012. "Deterministic Chaos and the Evolution of Meaning." *British Journal for the Philosophy of Science* 63:1–29.