

See discussions, stats, and author profiles for this publication at: <http://www.researchgate.net/publication/240819260>

Shannon – Boltzmann – Darwin: Redefining information (Part II)

ARTICLE *in* COGNITIVE SEMIOTICS · JANUARY 2008

DOI: 10.3726/81605_169

CITATIONS

4

DOWNLOADS

43

VIEWS

64

1 AUTHOR:



[Terrence W Deacon](#)

University of California, Berkeley

76 PUBLICATIONS 3,648 CITATIONS

SEE PROFILE

Shannon-Boltzmann-Darwin: redefining information. Part 2.

Terrence W. Deacon

University of California, Berkeley 2007

deacon@berkeley.edu

Abstract: A scientifically adequate theory of semiotic processes must ultimately be founded on a theory of information that can unify the physical, biological, cognitive, and computational uses of the concept. Unfortunately, no such unification exists, and more importantly, the causal status of informational content remains ambiguous as a result. Lacking this grounding, semiotic theories have tended to be predominantly phenomenological taxonomies rather than dynamical explanations of the representational processes of natural systems. This paper argues that the problem of information that prevents the development of a scientific semiotic theory is the necessity of analyzing it as a negative relationship: defined with respect to absence. This is cryptically implicit in concepts of design and function in biology, acknowledged in psychological and philosophical accounts of intentionality and content, and is explicitly formulated in the mathematical theory of communication (aka “information theory”). Beginning from the base established by Claude Shannon, which otherwise ignores issues of content, reference, and evaluation, this two part essay explores its relationship to two other higher-order theories that are also explicitly based on an analysis of absence: Boltzmann’s theory of thermodynamic entropy (in Part 1) and Darwin’s theory of natural selection (in Part 2). This comparison demonstrates that these theories are both formally homologous and hierarchically interdependent. Their synthesis into a general theory of entropy and information provides the necessary grounding for theories of function and semiosis.

Shannon, Boltzmann, and absence

Part 1 of this essay began by exploring the classic paradox of representational relationships: the fact that content is not an intrinsic property of whatever constitutes a sign or signal. The property of something that warrants calling it information in the usual sense is rather its linkage to something that the sign or signal is not. Classic conundrums about the nature of thought and meaning all trace their origin to this simple and obvious fact. This relationship has often been framed as a mapping or correspondence between a sign or idea in the mind and this something else, not present. This superficially reasonable account does not, however, distinguish the special nature of this relationship and what might distinguish it from other merely physical relationships. So, to use a classic example, the wax impression left by a signet ring is only wax, except for the mind that interprets it to represent the ring, the office of its bearer, and its bearer. But the wax impression is just wax and the ring is just a metallic form and their conjunction at a time when the wax was still warm and malleable was just a physical event, like so many other physical events where one object alters another when they are brought together. Something more makes the wax impression a sign that conveys information: it must be interpreted. Unfortunately, within this obvious answer a vicious regress hides. What we invoke with an interpreting mind is just what we hope to explain. The process we call interpretation is the generation of mental signs interpreting extrinsic signs, and we are left with the same problem inside as outside the mental world. The problem of specifying how a specific absent content inheres in some way in these components of the interpretive process is no better grounded in neurological processes than it is outside of brains.

As we saw in Part 1, a critical advance in developing a technically precise and practical conception of information came with Claude Shannon's introduction of a statistical mechanics approach to the analysis of signals and their capacity to carry information. Unfortunately, this analysis excluded any reference to problems of defining content or significance. But by bracketing these issues Shannon was able to provide an unambiguous, interpretation-free measure of what might be called the information-

bearing capacity. This analysis has stood the test of time with respect to any practical accounting of how much information a given medium can be expected to convey to its interpreter.

Although Shannon's conception of information totally ignores the issue of what information is *about*, or even that it is about anything, his analysis nevertheless provides an important clue for dealing with the absent content problem. This clue is provided by Shannon's negative characterization of information. Shannon's measure of the potential information conveyed by a given message received via a given communication channel is necessarily inseparable from the range of signals that could have been received, but were not. In other words, even though it ignores all issues of content and significance, Shannon information is also necessarily defined with respect to something absent.

In Part 1, it was further argued that the relevant physical property of the signal or sign medium is also a negative attribute. The reduction in signal entropy has the potential to carry information because it reflects the consequences of physical work and thus the openness of the signal medium to extrinsic influence. In thermodynamic terms, a change in the state of a physical system that would not otherwise occur, is inevitably characterized by a local reduction in its Boltzmann entropy (to distinguish it from Shannon entropy) resulting from work done on that system from outside. Since the information-bearing capacity of a signal is proportional to the improbability of its physical state, being in an improbable state likely reflects the effects of physical work. However, although this relation to work is fundamental, referential information can be conveyed both by the effect of work or evidence that no work has been done. Thus, no news can be news that something anticipated has not yet occurred. This demonstrates that Shannon information and referential information are not equivalent. This is because the signal constraint is not something located *in* the signal medium, it is rather a relationship between what is and what could have been its state at any given moment. A reduction in variability is a constraint, and a constraint is not an intrinsic property but a relational property. It is defined with respect to what is not present. So implicitly, a physical system that exhibits constraint is in that configuration due to extrinsic influences, but likewise if the sign medium exhibits no constraint or change from some stable state, it can be inferred that there was no extrinsic influence doing work on it. So the relationship of

present to absent forms of a sign medium embodies the openness of that medium to extrinsic intervention, whether or not it has occurred.

10) *The possibility of change due to work, not its actual effect, is the signal feature on which reference depends:* This is what allows absence itself, absence of change, or being in a highly probable state to be informative.

Consider, for example, a typo in a manuscript. It can be considered a reduction of information because it reflects a lapse in the constraint imposed by the language and necessary to convey the intended message, and yet it is also information about the proficiency of the typist, information that might be useful to a prospective employer. Or consider a technician diagnosing the nature of a video hardware problem by observing the way the image has become distorted. What is signal and what is noise is not intrinsic to the sign medium, because this is a determination with respect to reference. But in either case the deviation from a predicted or expected state is taken to refer to an otherwise unobserved cause. Similarly, a sign that doesn't exhibit the effects of extrinsic influence—for example, setting a burglar alarm to detect motion—can provide information that a possible event (a break-in) did not occur. In all cases, however, the referential capacity is dependent on physical work that has, or could have, altered the state of some medium open to extrinsic modification. This tells us that the link between Shannon entropy and Boltzmann entropy is not mere analogy or formal parallelism. This link is the ground of reference.

Interpretation preliminaries: making work and detecting error

Up to this point of the analysis it has been assumed that the relationships being described have involved signs and signals, and not merely physical events chosen at random. But in fact, *none* of the criteria specified thus far (including in Part 1) actually distinguish events and objects that convey information from those that don't. They are *requirements* for something to be information, but they do not in themselves constitute it. Most of what we mean by physical causality involves an account of physical work, and yet we don't usually describe just any physical change as constituting information. This dependence relationship between information and work has led to a promiscuity in the

use of the term ‘information’ that can be a source of confusion. Any physical difference *can* be interpreted as information about something else, whether it is the state of the mud on someone’s shoes or the presence and evenness of the microwave background radiation of the universe. Information is not something intrinsic to a sign, nor even just some relationship between a sign medium and extrinsic influences. It is something extrinsic, virtual, and possibly even discontinuous from any immediate causal history. And yet, as we’ve seen, it depends on the possibility of physical change.

Information, as it is understood commonly (as opposed to the more technical Shannon information), is a difference in something that is interpreted to refer to, or mean, something else for some interpretive purpose or other for some interpreter. This might suggest that at some point in our discussion we will need to introduce mentalistic concepts to cross the threshold into semiosis. Although this has been a common strategy—one that typically causes the natural sciences to part company from the semiotic sciences and humanities—the point of this analysis is to demonstrate that this is not a necessary presupposition. Recognizing that information is defined at every level of analysis—even with respect to signal properties (Shannon) and causal properties (Shannon-Boltzmann)—in terms of something absent, can help to bridge the gap between the merely physical and semiotic relationships that are being invoked.

The key to reframing this problem is to shift perspective from a structural to a dynamical approach. Although any physical difference can become significant and provide information about something else, this requires that—as Gregory Bateson quipped—it is a difference that *makes* a difference (Bateson, 1972). What does it mean to *make* a difference in this sense? Bateson’s aphorism turns on the ambiguity between two meanings of to “make” a difference: i.e. to matter and to cause to change. And it implies that the generation of information involves making a physical alteration in the state of things. As the previous analysis (above and Part 1) has shown, this requires performing work to drive conditions away from some more expected outcome; thus often away from intrinsic regularities.

Contrary to Bateson’s implication, however, physical work can also be described as a difference (e.g. an energy gradient such as an electric potential difference) that makes a difference (e.g. driving another system away from equilibrium). This parallel

between information and physical work is implicit in the Shannon-Boltzmann link described previously. But additionally, implicit in this aphorism is the idea that the difference that results is also work done as a consequence of the informing difference. So the difference in some sign medium conveys information because it induces some interpreting system to change its state and do work to alter something else. And specifically, in the case that the something else is connected in some way to the information.

The thermodynamic conception of work requires a non-equilibrium state, and the spontaneous tendency for this to progress toward equilibrium. Work results when a physical system in an unstable dynamic state is coupled to some other system and drives it further from its equilibrium condition. The recognition of a necessary linkage between concepts of information and work suggests that we also need to frame the discussion of the interpretive process in dynamical systems terms. I will argue that an interpretive process transforms a mere causal connection into a semiotic connection by virtue of how it organizes work in certain systems. To be capable of performing work, however, a system must be in a non-equilibrium state. But non-equilibrium conditions are inherently transient and self-undermining. This additionally means that any persistent non-equilibrium dynamics is necessarily dependent on something extrinsic to itself: a larger thermodynamic context from which the work to drive the system away from its most probable states is derived. Thus the presence of system in a non-equilibrium state entails environmental conditions that promoted it. Because non-equilibrium processes both presuppose work and are the necessary sources of future work, any full explanation of what constitutes an interpretive process must include a central role for non-equilibrium dynamics (see, for example, the classic work by Nicolis and Prigogine 1977).

Understanding that a process capable of generating information involves non-equilibrium dynamics provides a way to address the referential error problem, which has been a non-trivial problem for correspondence and mapping theories of reference, which has been around since the classic philosophical skepticism (e.g. Hume, 1739-40, Moore, 1903, and others). Following Bickhard (1998, 2000, 2002), I will argue, however, that the normativity behind representational error is an emergent property of the organization of

certain non–equilibrium processes irrespective of any component causal facts. This can lead to a very different appraisal of the problem.

A normative consideration requires comparison, and a means for defining and discerning between accurate and inaccurate reference. This comparison logic has its parallel in Claude Shannon’s (1949) analysis of the signal/noise problem in the transmission of a signal, even though Shannon’s theory is formulated irrespective of any specific referential relationship. Recall that Shannon demonstrated that any amount of unreliability in a communication process can be overcome by introducing a specified degree of redundancy into the signal, enabling an interpreter to utilize the correlations among similar components to distinguish signal from noise (so long as one can assume non-correlation between the noise source and the signal source). For any given degree of noise (signal error) below 100%, there is some level of redundant transmission and redundancy-checking, that can distinguish signal from noise. Effectively, the check for accuracy of transmission is self-consistency. But what is the check for accuracy of representation?

In one sense this problem can be seen as a higher order variant of the problem of correcting for signal noise, and it submits to a parallel solution. Thus, we often use the redundancy of interpretive consequences as a means for detecting representational error whenever this is available. In practical terms, this is the widely employed method of fact-checking. Comparing multiple independent reports of the same event can help to reduce interpretive error. For example, multiple witnesses to a crime who may have only observed some of the relevant events, and who may have poor memories of the details, or who may be withholding or falsifying evidence, can provide accounts that be compared and cross-checked to reconstruct the most probable course of events. Those accounts that have concordant reference are taken to provide the most likely and most accurate representations of what occurred.

It is also the essence of the method of the empirical sciences. So when an independent researcher replicates the results of another researcher’s experiments it reinforces confidence in the prior claim. The second researcher provisionally assumes the accuracy of the prior claims and operates accordingly. But whereas disconfirmation can lead to serious theoretical revisions, discovering consistency between results is only

the minimum. Using many independent methods, analogous to obtaining reports from multiple independent witnesses, and finding consistent results is even more convincing. This is why developing new tools for investigating the same phenomenon in different ways provides for a considerable increase in the representational confidence that is generated.

The logic of fact-checking differs in an important respect from introducing signal redundancy, however. This is because it is actually a means for increasing the entropy of the signal, not decreasing it, as would be the case for overcoming transmission noise. To detect representational error in this way, it is necessary to compare different and to some extent independent sources of information, and to instead take advantage of the otherwise uncorrelated signal diversity to overcome error. Each source of information will have its own idiosyncrasies to contribute, analogous to noise, but all will share, in common, being generated with respect to, and under the influence of, the same extrinsic events. So correcting representational error entails both an increase in the entropy of the signal—which increases the Shannon information—and taking advantage of redundancies in the constraints imposed on these processes by something extrinsic to each, but shared in common.

This is what might be called a second-order interpretation, since it treats each source of signal as an interpretive process, as well. In other words, it involves an assessment of relationships between alternative interpretive mechanisms, to the extent that each signal potentially reflects of some independent relationships between signal generation and the context of its production. Ultimately, then, this entails treating the constraints characterizing the different signal-generation systems as analogous to redundant messages themselves. Unfortunately, assuming a common source of redundant constraints in independent interpretations is never an infallible inference. This is because it involves an interpretation of similarity and difference, and there can be many reasons for not detecting difference, particularly when the Shannon entropy—instantiated by the number, complexity, and diversity of interpretive sources—is not large.

To explore this more carefully, consider again the example of the detective who compares many sources of information and uses their correlations to infer a common event, which they may or may not each indicate. Over time, as more interpretive

techniques have become available for this purpose—e.g DNA, materials analysis, and trace-elements detection—the interpretive redundancy increases, with an attendant increase in interpretive confidence. The detective’s problem, or that of a jury listening to a welter of potentially untrustworthy evidence, is to reduce the uncertainty of interpretation; i.e. to get at the “truth.” They must generate an interpretive response to the whole ensemble of sources of evidence and counter-evidence that best corresponds with what actually occurred beyond direct observation. The consistency (redundancy) and inconsistency (non-redundancy) of the evidence is not itself a guarantee that a given interpretation is accurate. Faced with the problem of comparing alternative interpretations of the same events, one is often forced to analyze other features of the source of the information to determine if there are systematic biases that might be introducing spurious or intentionally skewed levels of redundancy. Creating the false appearance of independent sources of information is, for example, a major tool employed in propaganda and confidence schemes.

Hypothesis-testing is most effective when the interpretive challenge involves objects or states of affairs that are immediately available for further exploration. In these cases the redundancy analysis can be fairly straight-forward. Hypothesis-testing involves behaving as though a given trial interpretation is accurate and observing the consequences of continuing to act in accordance with that interpretation to see if the consequences remain consistent with it. By acting in accord with a given interpretation, causal consequences of this interpretation can be generated to act as virtual new interpretations, each of which can be compared.

So, for example, on suspicion that a given business is corrupt, a law enforcement agency might set up a sting operation that will proceed as though their suspicion is true, and observe the consequences. The concept of producing actions that test for interpretive error, was hinted at in Bateson’s aphorism about information (cited above), and again involves the performance of work: acting to change circumstances to produce predictable results. All of these approaches to the problem of representational error-checking reinforce the claim that interpretation is a dynamical process that inevitably involves the generation of new information, in the form of new signals and new interactions that do work with respect to those which were generated previously. Although this shifts our

analysis upwards to a second and third order of information generation, the same core logic that we have seen at work in Shannon's classic analysis still applies: the information conveyed is determined with respect to the alternatives eliminated, whether about reliability of the signal, reliability of reference, or reliability of interpretation.

Darwinian information

In many respects this process of error-detection is crudely analogous the logic of natural selection, with an hypothesis as the analogue of a variant phenotype and the selective exclusion of certain of these based on their non-concordance with others as the analogue of selection. So an analysis of the logic of natural selection theory in terms of its parallels with this expansion of information theory may offer clues concerning how information can come to be reliably correlated with the physical workings of the world. Indeed, many theorists have compared scientific research and other truth-seeking enterprises to Darwinian processes (e.g. [Karl Popper's evolutionary epistemology](#), see for example Campbell 1974) and a number of contemporary philosophers have developed theories of function and mental content based roughly on the logic of natural selection (e.g., Fred Dretske, 1988 and Ruth Millikin, 1984). But a number of problems with these approaches have been uncovered, mostly having to do with information only being defined with respect to past conditions, not current conditions. , and there are questions about whether they can account for and detect error (see, for example, critiques by Bickhard 1998, 2000). To sort out these problems of the interpretive generation of information, then, we must first be sure we can answer the analogous problems posed by evolutionary theory.

In the standard Darwinian account of evolution by natural selection, many individual organisms with variant forms constitute a pool of options from which a small subset are able to successfully reproduce to generate the next generation. This subset succeeds because of their comparatively better fittedness to prevailing environmental conditions, and as a result of genetic inheritance the new pool of variant individuals that is produced inherits features from the parent generation that functioned best in that environment. By analogy to Shannon's model of the transmission of information, the initial variety of genotype and phenotype forms in the prior generation can be considered

the potential entropy of the lineage, and the reduction in “transmitted” forms that occurs due to differential reproduction and elimination processes therefore embodies information about the selective environment. With the evolution of each generation that is subject to selection there is therefore a decrease in the entropy of variant forms, (though there is also the incremental generation of new variations). Comparison of this reduced fraction of the variation of traits to the previous ancestral variety, including the non-reproduced variants, provides an assessment of the potential information generated in this transition. In theory, one should be able to quantify this entropy reduction for a given population of organisms for a given number of generations, and estimate the amount of Shannon information produced per time in the evolution of that lineage. It is this parallelism that warrants talking of evolution in informational terms, and ultimately for describing evolution as a process that produces new information.

11) *The measure of Darwinian information:* The reduction of genotypic and phenotypic variety by virtue of differential survival and reproduction of certain organism forms (favored in a given environment) is directly analogous to the reduction of signal entropy in Shannon’s analysis. Thus, in principle, evolved information can be quantified and rates of information evolution can be compared.

This reduction in “genetic” entropy is also evidence that outside influences have been involved. This is analogous to the way the reduction of Shannon entropy points to an outside influence constraining potential signal variety (e.g. a sender selecting a message or environment interacting with a scientific instrument) and thus exhibits the potential of physical work to alter the sign medium. In biological evolution, however, the outside source of influence is typically the environment within which this process takes place, and its role is essentially passive with respect to organism reproduction (the process that is analogous to signal generation in Shannon’s analysis). So the work being performed to “make a difference” in the generation of Darwinian information is not something extrinsic to the “living signal,” but rather is intrinsic to the self-maintaining and reproductive processes that characterize the organism. In other words, because the

non-equilibrium dynamics of life is selectively sensitive to its environmental context, the constraints implicit in the environment can become re-presented in the selective preservation of some living dynamics and not others.

Work is necessarily involved to generate information, but if the object of reference is a passive state, the source of this work must be intrinsic to the signal generation process itself. This organization is a characteristic shared by many scientific instruments that serve as detectors. By incessantly generating a far-from-equilibrium process, a device's intrinsic instabilities do work that can be used to exemplify their highly sensitive reactivity to certain contextual factors. A process that must continuously do work to maintain an unstable state requires specific conditions in its environment, and so its state can be used as an indication of the presence, absence, or change of these conditions. Often highly specific conditions are required for maintenance of unstable dynamics. Shifting away from these conditions can thus make a large difference because this potential to do work can amplify any tiny difference in some critical parameter into a large difference in the dynamics of the signal medium.

Though highly specific sensitivity reduces the diversity of what can thereby be conveyed by a change in this dynamic signal, it can also provide exquisite precision of reference. The result is an effective increase in the entropy of the referential (Shannon-Boltzmann) information, even if the signal (Shannon) entropy is low; e.g. perhaps only a yes/no indication of crossing some threshold value. Greater precision of measurement effectively multiplies the potential entropy implicit in the different states of the represented object, even if only one critical threshold change of value is being signaled. This is intuitively reasonable since the greater the certainty, the less the uncertainty; and reduction of uncertainty is the measure of Shannon information. The canary in the mine can thus tell the miner that although he is not yet gasping for air, that possibility is incrementally close. Similarly, a metal detector is highly sensitive to the presence or absence of an object capable of being attracted to a magnet, but little else, because only conductive metals can disrupt the detector's magnetic field. Despite the low entropy of the signal, this specificity is what the treasure-hunter or mine sweeper wants: vastly reduced representational uncertainty (precision of reference). While this is intuitively obvious, it often gets ignored in technical discussions that do not distinguish these two

levels of information. It again demonstrates the non-correlation between assessments of information and entropy at these different levels of analysis:

11) *Referential entropy reduction does not necessary correlate with Shannon entropy reduction:* Referential entropy is a function of the possible states of the object of reference, even if most of these potential differences do not correspond to specified differences of a coupled signal-generating medium.

The above examples still appeal to an outside observer to interpret them as information; life does not. Living organisms are far-from-equilibrium systems that are dependent on transforming the physics of their interaction with the environment into information about that environment. They also pass information about the relationship between their internal dynamics and the environment to succeeding generations. They do this primarily via genetic inheritance, but also through cytoplasmic overlap, social interaction, and environmental modification. And this inherited information is “interpreted” in the form of organism structures and processes that compensate for variations in that relationship which tend to perturb the stability of these organism features.

This *intrinsic* interpretive capacity derives from an additional and crucial feature that characterizes living processes over and above their far-from-equilibrium thermodynamics: production and maintenance of the mechanisms and processes that maintain this dynamic organization with respect to environmental fluctuations. So if we treat reproduction as analogous to signal production, then the material, dynamical, and structural match between these organism processes and conditions of the local environment are the determinants of the evolved information. This is why organism features reflect the constraints of their environments, and what justifies claiming that the information embodied in a phenotype and its genotype is information *about* this environment. But this *aboutness* only reflects these environmental constraints *with respect to* the constraints of the formative processes critical to reproduction and self-maintenance of that organism.

Both adaptations and functions are defined with respect to something they are not, something other, and are often in service of something not yet in existence. Adaptations are defined with respect to something about the world outside the organism that significantly affects the organism's means of maintaining and reproducing these adaptations.

12) *The information of adaptation:* Adaptation is defined with respect to the constraints on phenotypic variety that embody constraints of the environment which are both presupposed by its existence and yet extrinsic to its far-from-equilibrium processes.

The signal/noise distinction is not intrinsic

There is one further critically important difference between the abstract logic of communication theory and the evolutionary process: the shifting status of signal versus noise in evolution. If we liken the transmission of traits from generation to generation via reproduction to signal transmission over a communication channel, then mutation in biology is the analogue of noise introduced into a communication channel. In most communication processes noise is a nuisance. It degrades the information by introducing new uncorrelated entropy into the signal, and this increases uncertainty about what is signal and what is not, thereby potentially corrupting the message. But whereas the introduction of noise decreases the potential Shannon information capacity of a channel, it paradoxically increases the capacity for reference, because it increases total Shannon entropy. It is as though an additional information channel is available, because noise is also a consequence of the openness of the physical system that is being used as a sign medium, and so it too reflects some source of signal modification besides that which the sender provides. Of course, noise is just noise if you are only interested in what was originally sent and not interested in the cause of the degradation of that signal. And yet, this normative decision depends on the interpretation process. Noise can be signal to a repairman.

From the perspective of Shannon information, noise is a source of equivocation or ambiguity in the signal. A noisy signal, like a text containing typos, contains signals

replaced by uncorrelated alternatives. Shannon's analysis showed that it is possible to compensate for equivocation between signal and non-signal if the transmission and interpretation processes can take advantage of signal constraints or redundancies. In the evolution of adaptive phenotypes, however, there is no such shared expectation to go on. Understanding how this is accomplished when there is no context of introduced redundancy to rely upon is the critical clue to explaining how evolution ultimately transforms noise into signal.

So what if there is no information available in the signal to help discern transmitted from randomly substituted bits? Consider the case of a set of instructions in which there are word substitution errors, but in which there is no violation of meaning, spelling, or grammar to indicate that it is inappropriate. This sometimes happens with foreign-made devices that come with assembly or use instructions that have been poorly translated from an unfamiliar language. In these circumstances we often provisionally assume that the instructions are accurate and attempt to accomplish the task described. If in the process we find that something doesn't work out as described we may suspect error in the instructions, and careful attention to the task described can often provide clues to the locus of this error.

This trial-and-error approach is also a form of hypothesis testing, as described above. The redundancy being relied upon is between the referential information in the communication, and the constraints of the application context. If the information accurately represents features of some physical system (for example, the instructions about operating some mechanical device) its interpretation in terms of the actions performed on (or interactions with) that system will correlate well with physical constraints required to achieve a given expected result.

- 13) *Representational error-detection via predictive interaction*: The reference of a sign or signal is also susceptible to error-correction via redundancy at a higher level than signal organization. The reference of a signal implicitly makes a prediction about certain extrinsic causal possibilities. Physical or logical interactions with these extrinsic conditions will be constrained to either conform

or not conform with this prediction, and if not this will disconfirm the represented state.

So the logic of natural selection is analogous in many ways to a trial and error process, except that in natural selection there is no extrinsic source of representation to check against. Success or failure to reproduce is all that distinguishes representational accuracy of the information embodied in the genotype and phenotype. But reproduction allows for further iterative testing of these interpretive consequences. So the succeeding generations effectively stand in for the outside source of comparison necessary for error-checking. If one's genetic inheritance contributes to producing a body with appropriate adaptations, it is because the constraints it embodies are in some degree of correspondence with constraints of the environment. Unfortunately, dead men tell no tales, as the cruel aphorism suggests. So it would seem that there is no recording of the many failures-to-correspond; no independent representation of which were the errors and which were not.

What counts as useful information in biological evolution is determined *after the fact* with respect to its ability to pass through the functional error-correction mechanism of natural selection. The “accuracy” of the inheritance signal is both “tested” and refined by the way the far-from-equilibrium dynamics that constitutes organism development, maintenance, and reproduction conform to environmental constraints and opportunities. Evolution is thus a generator of information for the organism and a process that rectifies this information with its reference. Thus, although the evolutionary process is itself non-normative, it produces organisms capable of making normative assessments of the information they receive.

- 14) *Evolution generates and rectifies referential information:* Interpretive processes do work and therefore involve non-equilibrium processes that are necessarily context-sensitive. The entropy of an ensemble of interpretive processes arising from uncorrelated sources and thus lacking prior reference (‘noise’ from the perspective of its source) can come to acquire reference because the constraints

that are serendipitously embodied by it happen to also be redundant with certain boundary conditions supporting this dynamics.

15) *Pragmatic convergence*: The evolutionary process can progressively increase the functional correspondence between organism dynamics and contextual preconditions. To the extent that the constraints of this dynamics conform to environmental constraints that are consistent with its continuation, these intrinsic constraints embody this correspondence in the ongoing dynamics. Although this does not constitute information, this process is the basis of referential rectification.

Students of evolution have not usually insisted that the absence of the lineages that go extinct is what determines the functionality of the traits that persist. One could see the surviving lineages and their adaptations through the lens of engineering design in terms of identified functions that were designed to achieve a previously specified purpose. But although this analogy has a superficial attractiveness, it is undermined by the fact that few if any biological structures can be said to have only one distinguishing function. Their fittedness, internally and externally, is irreducibly systemic because adaptations are the remainders of a larger cohort of variants that are selected with respect to one another and their environmental context. There is no simple mapping of genetic-phenotypic information and adaptive function. So that which constitutes the reference of the inherited information is ultimately defined only negatively (i.e. by constraint). Biological function is not, then, positively constructed but is rather the evolutionary remainder that occupies the constrained space of functional correlations that have not been eliminated. This is the basis for novel functions to emerge in evolution as well as the possibility for evolutionary *exaptation* (the shift from one adaptive function to another). In this respect, genetic information is neither merely retrospective—i.e. about successful adaptation in the past—nor does it anticipate future novel adaptations. It is not an aspect of a static relationship, but emerges in process, as its interpretive consequences perform work that may or may not turn out to support this process continuing.

The emergent nature of information

Although the account so far has been framed in terms of the information involved in biological evolution, this model is generalizable to other domains. The nested dependencies of the three levels of entropy-reduction—here characterized by Shannon’s, Boltzmann’s, and Darwin’s variations on this theme of entropy-reduction—define a recursive architecture that demonstrates three hierarchically nested notions of information. These three very roughly parallel the classic hierarchic distinctions between syntax (Shannon), semantics (add Boltzmann), and pragmatics (add Darwin). They also roughly parallel the relationship between data, content, and significance, though to understand how these semiotic levels are interrelated we must carry this analysis out of the realm of biology and into the domain of communication.

The appeal to Darwinian selection as the ultimate mechanism for the generation of new information relationships might suggest that we should take a strictly etiological view of information. In other words, we might be tempted to argue that information is only discernable *post hoc*, after selection, and after entropy reduction. But this is misleading both for biology and for information relationships in general. As with biological function, the specific selection history of a given representational capacity may be necessary to explain present usage, but past correspondences are not what it is currently about. Past correspondences have improved the chances for reliable and precise predictive correspondence, but it is a relationship to the present condition that matters. The very fact that information and noise are not intrinsically distinguished, and that mutational noise can become biological information in the course of evolution, exemplifies this property.

This is the problem with simple etiological explanations of adaptive function and representation, which treat information and function as retrospectively determined by their selection history. Because information-generating processes emerge in systems constituted by a pragmatic selection history, the ground of the correspondence between information and context is determined negatively, so to speak, by virtue of possible correspondences that have been eliminated, but it leaves open the issue of correspondences never presented. No *specific* correspondence is embodied with full precision and present correspondence is not guaranteed. With functional correspondence

under-determined, novel functions can arise *de novo* in unprecedented contexts and incidental properties of the sign or signal may come to serendipitously serve emergent functions. In short, while the possibility of information generation and interpretation depends on a specific physical selection history, the present influence of this information on the persistence of the system that enables it may be serendipitously unrelated to this history. This is the basis for the evolution of new function but it is also why information is always potentially fallible.

The evolutionary process is not, however, a normative process. Conditions can be good or bad for an organism, or for life in general; an organism's responses to the world can be effective or ineffective in achieving its intrinsic ends, and its adaptational dynamics can accurately or inaccurately link changes in organism activity with changes in extrinsic conditions contributing to the persistence of this dynamical organization; but evolution just occurs. So although the evolutionary process can further the pragmatic convergence between interpreted content and extrinsic reference, information is not in any sense available *to* evolution, only to the organisms that are its products. Evolution generates the capacity to interpret something *as* information. This capacity is intrinsic to a self-perpetuating far-from-equilibrium system that depends on its environment and does work to modify that environment in a way that reinforces its persistence. Information is a relational property defined with respect to this persistently unstable dynamical regularity, or as the philosopher Charles Sander Peirce (1931-35) would have said, with respect to a "habit"—understood in its most generic sense; specifically, a self-perpetuating self-rectifying habit.

So genetic information is about cellular chemical reaction possibilities, their roles in constituting the organism, and how this relationship between genes and their effects also correlated with extrinsic conditions that supported the maintenance of these possibilities in the past. It is information about organism design and function because it introduces critical constraints into the non-equilibrium processes that may ultimately contribute to the perpetuation of that relationship. It is interpreted by the persistence of the self-perpetuating process that it contributes to. It is not information about the present world, only extrinsic signs can fill this role. The ability to make use of environmental

features as information nevertheless depends upon this closed interpretation of genetic information for the ability to obtain information from extrinsic sources.

This ability to use extrinsically generated events and objects as information derives from the special dynamics of living processes. Because organisms are constituted by specially organized persistent far-from-equilibrium processes, they are intrinsically incomplete. In this regard they are processes organized around absence. Not only are biological adaptations evolved and defined with respect to features of the world extrinsic to the organism, but in many respects these are only potential features which are also absent from the current environment. Thus, adaptations of an organism that have to deal with unusual conditions, like high altitude or extremes of heat or cold, may never be expressed in a lifetime. For this reason the maintenance of intrinsically unstable far-from-equilibrium conditions entails mechanisms that effectively anticipate the possible variations of environmental conditions. But they do so with respect to a living process that is at the same time incessantly asymmetrically directed contrary to high probability states in multiple ways: they do work (a) to maintain their far-from-equilibrium state (which supports persistence of the ability to do work), (b) to generate specific organic forms (i.e. they constrain dynamical processes and generate structures which have highly constrained low probability features), and (c) to achieve the specific outcome of maintaining themselves long enough to reproduce the global organization supporting processes a, b, and c. So with respect to these three improbably asymmetric dynamics there are many critical extrinsic factors that are relevant. This combination of absence and necessary relevance to an asymmetric process, incessantly interacting with and modifying the world, is what projects the property of information into otherwise merely physical states and events.

Consider a non-mentalistic example: deciduous tree which alters its metabolism in response to decreasing day length and cooling temperatures in the early months of Autumn, resulting in the eventual withdrawal of metabolic support for its leaves so that they dry up and eventually become severed from the branches they grew from. This adaptation to the difficulties of winter involves a mechanism that treats these environmental changes as information about likely future events that would have an impact on survival and effective reproduction. Insofar as this response has, in previous

generations, resulted in persistence of the lineage compared to others lacking it, the mechanism has acquired interpretive reliability. The reliability of the seasonal changes in these factors provides constrained variation to which the constraints of the tree's metabolic mechanisms have become tuned. But it is not merely these correlations that constitute the informational property of these seasonal changes for the tree. The day length and mean temperatures are also correlated, but one is not intrinsically information about the other. It is only with respect to the end-directed improbable dynamics of the tree's metabolic processes that one or the other of these is informative; and specifically informative about boundary conditions potentially affecting that dynamics.

At one point I worked in an office near a number of trees of the same species that had been planted as part of the landscape design for the campus. A few of these trees, which were planted close to an automated streetlamp and next to the exhaust from the building's ventilation system, always were very late to change the color of their leaves and drop them, compared to the others. On the one hand, one might argue that these few trees were misinterpreting these artificial signs, because they don't accurately represent seasonal changes. On the other hand, to the extent that these artificial conditions were nevertheless reliably predictive of local factors affecting the trees' metabolism, one would be justified in arguing that the interpretation was correct, because it promoted the dynamical outcome by virtue of which the mechanism exists. This shifts the focus from the evolved function to the immediate incremental consequence of the evolved mechanism as the ground for referential information. The evolved mechanism constrains the dynamics of possible interpretation, but doesn't determine it. Each moment of interpretation is in some way supportive or disruptive of the self-maintenance of this dynamical trend. This means that not only is there an historical origin for the normative property of this interpretive process, there is also an ahistorical and immediately efficacious normative property as well. And this need not be consistent with its evolved function. In fact, this possibility is a necessary condition for evolution, since essentially every adaptation has evolved from prior forms and mechanisms that often served very different adaptive functions (such as feathers originally evolving as a form of insulation).

Function and representation are made possible by the way living processes are intrinsically organized around absent and extrinsic factors, and the Darwinian process

inevitably generates increasingly convoluted forms of dependency on absence. Information is a relational property that emerges from nested layers of constraint: constraints of signal probability (Shannon), constraints of the dynamics of signal generation (Boltzmann), and the constraints required for self-maintaining far-from-equilibrium end-directed dynamics (Darwin). Because information is a relationship among levels of constraint generated by intrinsically unstable physical processes, it is also normative with respect to those processes. But constraint is a negative property, and thus neither something intrinsic nor determinate. This means it is intrinsically incomplete and fallible. Yet it is these very properties that make it evolvable, and though relative, potentially and indefinitely refinable.

Information in language

This analysis of the concept of information addresses only the most basic problems relevant to a theory of representation and semiotic processes in general. Many more issues need to be addressed before we can apply some of these insights to problems of our mental lives. But even just this generic analysis can provide a useful perspective to reflect on issues of interest for the study of communication and language. For example, we tend to think of word-reference in positive terms, i.e. as a correspondence relationship between a term and some concept and between both and a selected set of objects, events, or properties of things in the world. Arguments in the field have for this reason often focused on trying to define the nature of this correspondence, the problem of locating or specifying the ontology of the ‘content’ of information, or determining the status of the objects of reference (e.g. whether a class, a general concept, or individuals). But many of these issues can be usefully reframed in constraint terms.

This is well illustrated by the metaphoric capacity of words to take on novel functional roles and to generate new meanings and references. Consider a word like ‘shadow,’ which has an unambiguous meaning and refers to a commonly experienced phenomenon. Owing to the combinatorial power of language the following metaphoric extensions have become commonplace: a shadow of suspicion, a rain shadow, a shadow of doubt, living in the shadow of one’s father, an apprentice shadowing his mentor, having a shadowy past, and so forth (see also Hofstadter 2001). Such uses are not

particularly dependent on poetic sensibilities, which tend to trade on far more obscure and indirect metaphoric inferences, and are easily understood the first time they are heard. This capacity is often described in terms of a finite set of abstract semantic features whose presence or absence is invoked whenever the term is used; for example physical, inanimate, singular, nonmaterial, etc. Thus these features are combined with those evoked by the other terms in the various compounds. But are these features in any fundamental way different from the corresponding descriptive words for them? It might be complained that this analysis merely shifts the problem to explaining the semantics of these other “features,” understood as only slightly more basic quasi-words for concepts that are a bit easier to treat as simply present or absent. And indeed, the same kind of metaphoric examples could be compounded using the terms for these features. Doesn’t this merely beg the question, leaving the problem of explaining this informational relationship untouched? But here is where considering this correspondence as analogous to biological adaptation and exaptation may be helpful.

Although for both organisms and languages the history of past functional successes is constituted by specific instances of correlation with specific contexts, this specificity is neither a necessary nor sufficient condition for current functionality or current representation. The immediate context of interpretation and use are likely unique, only bearing a family resemblance to any previous instance that has conditioned the current use of this information. What ultimately matters is the outcome of this immediate process with respect to the continued coherence of that process. With respect to biological traits the outcome will matter if it aids or diminishes the perpetuation of that trait or habit of use, and thereby contribute to the system that this depends on. With respect to language communication the outcome will matter to the extent that it accomplishes some personal goal, but also to the extent that this affects the habit of expressing the information (description, request, demand, etc.) in that particular way. And across vast numbers of uses over many generations the cumulative usefulness, appropriateness, and catchiness of a particular form of expression will likely also influence its probability of recurrence. The specific past history results in some habits of expression (of traits or word uses) being dropped, others being retained, and many possibilities never being realized. But what is not eliminated is left available, and cannot

then determine a neat one-to-one mapping of adaptation to environment, function to use, word to reference; only a fuzzy correlation between two constrained and under-explored realms. In an uncertain world, each instance is slightly different, each context is in some measure novel, and—because of this negatively determined correlation—each use is both susceptible to failure and to discovering novel functionality. So whereas evolution and etiology *explain* the persistence of an organism trait or a linguistic usage, respectively, these historical origins do not *determine* either function or reference, but merely bias the likelihood of exhibiting some correspondence that is functional.

From information to semiosis

Information is a property (or more accurately a three-tiered set of interdependent properties) of any semiotic process. As we have seen, the common denominator of all three properties is constraint with respect to a possible entropy or range of variations. In Peircean terms (Peirce 1931-35), Shannon information is prescinded from Shannon-Boltzmann information which is prescinded from Shannon-Boltzmann-Darwin information. The two prescinded analyses thus assume a process conception, but bracket it from consideration to focus on successively more limited component constraint relationships. But it is not merely process that is the critical feature, but rather a particular topology of process characterized by persistent self-maintaining self-reconstituting far-from-equilibrium dynamics. This special class of recursively constrained dynamics thus confers a normative character to its component relationships because of its intrinsic dynamical asymmetry. In addition, it is a dynamics that effectively transforms the causal constraint of extrinsic physical relationships into constraints on the incompleteness of the far-from-equilibrium process. Elsewhere (e.g. Deacon 2006) I have described this dynamical incompleteness—characteristic of both living adaptations and of representations—as a “constitutive absence,” because the constraints that define this topology are literally constituted by something extrinsic and not included in the components or their relationships. It is this constraint logic that is the defining characteristic of semiosis. But it is also the logic that links the properties of serving a function and conveying information to a specific class of physical processes, by virtue of the absent options and reduced degrees of freedom that characterize them.

The Darwinian analysis largely focuses on function, which is well accounted for in biological terms, but with respect to language this is clearly not sufficient to account for the kind of representation that words provide, and so too for many other forms of human-scale semiosis. The dynamical analysis provided here can be seen as providing the functional ground out of which all semiosis must emerge, but although it explains what might be called the extended physics of semiosis, there is a higher-order analysis left unspecified that involves extrapolating from the information relationship to ways this can be differently used in thought and communication. Part of this difference is exemplified by the fact that there are different modes of representational relationships (e.g. iconic, indexical, and symbolic) that depend on different relationships between the features of the sign/signal medium and what they provide information about. This broader realm of semiotic relationships is a much larger domain than merely concerns informational relationships. In many ways, both iconic and symbolic relationships are less and more than informational, respectively, and only indexical relationships directly provide information. One might characterize iconic relationships as presenting the possibility of being used to acquire information, though not being a source of information, and one might characterize symbolic relationships as exemplifying relationships between forms of information.

Does this mean that we must add a fourth step to this analysis of information beyond the contributions of Shannon, Boltzmann, and Darwin? Not exactly. There is an important parallel between this dynamical hierarchy and the icon, index, symbol hierarchy that can point the way to a general theory of semiosis that fully integrates the physical-informational analysis with the representational analysis. Indeed, both exemplify Peirce's categorical hierarchy of prescinded relationships: Firstness, Secondness and Thirdness; one in dynamical terms and the other in formal-correspondence terms. Although to fully unpack this parallel logic and complete the process of unifying dynamical and semiotic theories will require an analysis that is far beyond the scope of this essay, the information question can itself be prescinded from the broader challenge of formulating a thoroughly naturalized conception of semiosis. Since it is to some extent prior to and presumed within semiotic analyses, clarifying the information-dynamics relationship is a necessary first step.

In our Darwinian framework, how something was successfully interpreted in the past into dynamical continuation becomes a source of constraints embodied in an interpretive habit generated in the present and with respect to its own extrinsic boundary conditions. The entropy embodied in the constraints of this inherited-persisting sign/signal-generating process determines what fraction of the environmental entropy will be re-presentable in this transition from medium to medium. So in this respect the interpretation of any new sign/signal is inevitably constrained (at a higher level) by the functional consequences of past interpretations and the relationship between these will in turn be a constraint interpreted by the functional consequence of aiding or impeding the preservation of this habit. The information embodied in remembered (or inherited) interpretation-consequence relationships is what picks out certain constraints of a current potential source of information by virtue of the signal entropy it thereby ignores. Whether this particular selection of signal entropy picks out the appropriate or inappropriate information is a function of the support it provides to the persistence of the system within which this habit of interpretation has developed. And this is determined by the higher level selection imposed by the constraints of the environment on this system's requirements for persistence.

While this ultimate interpretive level occurs external to and between organisms in the biological world, in the analysis of mental processes this same dynamical logic must be recapitulated in brains. This is a Shannon-Boltzmann-Darwin dynamic predicated on a lower-level Shannon-Boltzmann-Darwin process. This additional convolution of the logic embeds the representational information generating dynamic in a functional information generating dynamic. This is what makes the semiosis of mind significantly more complex and subtle than the semiosis of biological function alone, and it also makes semiosis as understood mentalistically parasitic on semiosis understood functionally.

This representational aspect of the information is thus doubly open to further interpretive possibility. Not only can function be an emergent outcome of evolution, but within the higher-order context of representational relationships embedded in and dependent on functional relationships, new representations can emerge with respect to emergent functions.

Emergent meanings and unprecedented referential relationships are constantly being generated in everyday language usage, some like the metaphoric extensions noted above, but others like the technical reuse of the term ‘energy’ that has now become a ubiquitous fixture in modern folk physics. The penumbra of representational possibilities that is left available is a critical prerequisite for the emergence of new meanings and the adaptability of language. Much of this generativity is accounted for by the incredible combinatorial use of prior representations, but even this depends on the openness of referential possibility implicit in its function-based foundation. So attempts to coin new terms (such as those introduced in this book in order to designate concepts for which existing terms are unavailable) succeed best if they borrow meanings, functions, and connotations from other words or morphemes (e.g. from ancestral languages such as ancient Greek) and thus take advantage of the undifferentiated possibilities that they embody. To invoke a term coined by the psychologist James Gibson, the constraints generated by this selection history don’t pre-determine possible uses, they instead create *affordances*.

Probably the most significant contribution of this analysis of information is not, however, this contribution to the understanding of functional and representational openness, but rather the way that it explicitly identifies the nestedness of information within physical processes, and thus its relationship to concepts of energy and form. This approach bridges the classic Cartesian gulf between the extended world of physical interactions and the presumed non-extended world of mind and representations, because it both recognizes the essential physicality of information and yet demonstrates that what is informative in a sign or signal is something precisely not present; and thus without extension. At every level, the emergence of informational relationships is the result of the constraints exemplified by via difference and produced by various pragmatic reduction processes. Information, reference, and fittedness are each differentiated with respect to something absent, and these relationships are progressively nested within processes similarly organized at higher levels. From this perspective Cartesian dualism can be traced to a failure to understand the causal significance of absence and constraint in physical processes, leading to the identification of Mind with a disembodied realm, and a paradoxical conception of mental causality.

Recognizing the centrality of this reduction-absence logic is particularly instructive for demonstrating how easy it is to become seduced by what we might call the fallacy of “misplaced aboutness.” This is implicit in eliminativist alternatives to Cartesian dualism and it is ubiquitous in computational conceptions of evolution, mind, language, and social processes. It is an analogue to the Whiteheadian fallacy of misplaced concreteness, and involves the tendency to identify the content or ground of an intentional relationship with some one or more of the substrates of the process; e.g. in a quantum bit, a computer algorithm, a DNA molecule, a neural circuit, a word-object correlation, and so on. What the present analysis purports to demonstrate is that the representational relationship cannot be vested in any object or structure or sign vehicle, it is not reducible to any specific physical distinction, nor is it fully constituted by a correspondence relationship. But neither is it a primitive unanalyzable property of minds. Instead, even simple functional and representational relationships emerge from a nested interdependence of generative processes that are distinctive only in so far as they embody specific absences in their dynamics and their relationships to one another. These absences embody, *in the negative*, the constraints imposed on the physical substrates of signals, thoughts, and communications that can be transferred from one substrate to another, and which thereby play efficacious roles in the world as inherited constraints on what tends to occur, rather than acting as pushes or pulls forcing events in one direction or another. Constraints don’t *do* work, but they are the scaffolding on which the capacity to do work depends.

Conclusions

This is only the barest outline of an information theory that is sufficient to account for some of the most basic features of functional and representational relationships, so it cannot be expected to span the entire gap from biological function to conscious agency. But considering that even very elementary accounts of biological function and representation are currently little more than analogies to man-made machines and human communications, even a general schema that offers a constructive rather than a merely descriptive analogical approach is an important advance. It shows that information theoretic and semiotic approaches to function and cognition are not incompatible. But

until this schematic analysis can be formalized in more precise mathematical terms, many questions about the shape of such a synthesis will remain unanswered. And even were this to be completed, it would still be far from adequate to address the range of issues relevant to the cognitive sciences, since many levels of complex self-organizing and selection processes—and thus information-generating processes—intervene between the simplest forms of functional adaptation and human symbolically mediated intentionality. Spanning these gaps will require considerable empirical work informed by this sort of synthetic analysis.

Nevertheless, despite its schematic and merely descriptive form, this exploration of the relationship between information theory, thermodynamics, and natural selection unpacks some of the unrecognized complexity hidden within the concept of information. By generalizing the insight captured by Claude Shannon’s equation of information with entropy-reduction and tracing its linkage to analogues in thermodynamic and evolutionary domains we have been able to address a few of the most vexing issues of representation, reference, and normativity (i.e. usefulness). These inadequacies in current definitions of information have posed seemingly insurmountable obstacles to formulating a theory of representation sufficiently rich to serve as the basis for empirical science, sufficiently subtle to explain the curious “inexistence” of content in the signs and signals that convey it (to use Frans Brentano’s term for this special sort of absence), and sufficiently grounded in physics to explain representational fallibility, error-checking, information creation, and the relationship between informational and energetic processes. By providing an account that expands traditional information theory to address these issues of physical efficacy, representation, and normativity, this argument provides a necessary first step toward a scientific semiotic theory.

References:

- Bateson, Gregory (1972) *Steps to an ecology of mind*. New York: Ballantine Books.
- Bickhard, Mark H. (1998) Levels of Representationality. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(2), 179-215.
- Bickhard, Mark H. (2000) Autonomy, Function, and Representation. *Communication and*

- Cognition – Artificial Intelligence*, 17(3-4), 111-131.
- Bickhard, Mark H. (2002) The Biological Emergence of Representation. In T. Brown, L. Smith (Ed.) *Emergence and Reduction: Proceedings of the 29th Annual Symposium of the Jean Piaget Society*. Erlbaum, pp. 105-131.
- Deacon, Terrence W. (2006) Emergence: The Hole at the Wheel's Hub. In P. Clayton & P. Davies (eds.) *The Re-Emergence of Emergence*. Cambridge, MA: MIT Press, pp. 111-150.
- Dretske, Fred I. (1988) *Explaining Behavior*. Cambridge, MA: MIT Press.
- Hofstadter, Douglas (2001) Analogy as the Core of Cognition. In *The Analogical Mind: Perspectives from Cognitive Science*, edited by Dedre Gentner, Keith J. Holyoak, and Boicho N. Kokinov, Cambridge MA: The MIT Press/Bradford Book, pp. 499-538.
- Hume, David (1978) *A Treatise of Human Nature*. Index by L. A. Selby-Bigge; Notes by P. H. Nidditch. Oxford.
- Millikan, Ruth G. (1984) *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Moore, George E. (1903). *Principia Ethica*. Cambridge: Cambridge University Press.
- Nicolis, G., Prigogine, I. (1977) *Self-Organization in Nonequilibrium Systems*. New York: Wiley.
- Peirce, Charles S. (1931-35) *Collected Papers, Vol. 1-6*, edited by Charles Hartshorne and Paul Weiss. Cambridge, MA: Harvard University Press.