

Consequences of a functional account of information

Stephen Mann

August 2017 draft. Published version at *Review of Philosophy and Psychology*.

1 Introduction

Several misconceptions about the application of information theory in natural science are widespread in philosophy. It is therefore necessary to clarify the assumptions required for informational measurements to have significance. Contemporary accounts of informational measurements in the natural sciences are stuck between “the Scylla of meaning and the Charybdis of causation.” [26, p. 191]. Here I promote a view that retains the benefits of each while dispensing with problematic commitments. Information is a contribution to accurate performance of a function. It increases the efficiency of function performance by better apportioning the physical resources available.

In the next section an outline of the logical map of information concepts is presented. Assumptions in various domains lead to differing significance for different information measurements. A lightweight understanding of the positive thesis – what a functional account of information looks like *roughly* – follows. Three advantages of a functional approach are then discussed. First, traditional emphasis on the indicative aspect of information is supplemented with its counterpart instructional aspect, supporting contemporary work on subpersonal content. Second, an influential distinction between causal and semantic information, argued recently to be misconceived, is further denigrated. Finally, mathematical results adverting to utility are recast in terms of function, allowing us to apply information theory in evolutionary biology, neuroscience and elsewhere.

2 Perspectives on information theory

In order to properly locate the positive thesis, an overview of recent philosophical claims about the application of information theory across the natural sciences is in order.

Phenomena in several subdisciplines of biology and cognitive science recommend the use of information theoretic formalism. But differing assumptions entail different interpretations of formal results. Unfortunately, the special details of mathematical communication theory lead many to import one of its central tenets into areas where it does not belong. It is often claimed that Shannon and Weaver established that information theoretic formalism, in any domain, is irrelevant for the meaning of information transmitted. The claim is false but widely believed [7, p. 344 col. 2] [21, p. 395] [38, p. 21] [25, p. 1984]. Philosophical understanding of information in natural science is misshapen. The problem is no doubt caused by many factors, but one particular fallacy stands out.

Mathematical communication theory (MCT) deals with symbols that stand for symbols. The coded message for which information is quantified represents a primary message whose meaning is irrelevant for this quantification. In other contexts, information can be quantified for symbols that stand for things other than symbols. The engineering context is special, but its mathematics are general. A coherent account of information, and an understanding of the fruitful application of existing mathematical tools across natural science, results from rejecting the assumption that only meta-symbols can be quantified. This is the subject of the present essay.

The remainder of this section introduces MCT, after which section 3 promotes a generalisation of the concept of information that renders it suitable for application in various domains of natural science. Three advantages of this approach are then discussed. First, a new perspective on the coded messages of MCT is described in section 4. The distinction between natural and intentional information is explored in section 5. Finally, section 6 examines recent applications of information theory in microbiology and cognitive science. I argue these treatments accord well with a functional account.

Two key components of mathematical communication theory are pressed into service in philosophy. Both were introduced by Shannon as part of his foundational text. First, the central transmitter-receiver model finds analogies across science, and philosophy helps determine whether purported applications are appropriate. Second, the engineering problem whose solution is the aim of communication theory has sometimes been taken as fundamental to the broader domain of information theory. The crux of the present essay is that contemporary philosophical interpretations of both of these factors are largely incorrect. In the remainder of this section I survey them in order.

2.1 The central model

In MCT, information is quantified as the extent to which the source message can be recovered at the target. It is a function of statistical properties of the source message and the channel through which it is sent, and as such is nothing to do with its meaning (insofar as statistical properties of a source have nothing to do with meaning).

As an example of communication, the central model is a rather peculiar case. Rarely does it apply exactly outside communications engineering. In order to apply information theory in other domains, the formalism has to be generalised. To this end, consider one of the special properties of messages in the central model. Because the symbols are about symbols, encoded messages usually carry two meanings, one folded within the other. The first is the instruction how to recover the original message.¹ The second is the meaning of the original message. It is often pointed out that the formalism of information theory is blind to the second meaning. But there are communication systems whose messages are not encoded in this way. In these systems, the meanings of messages are precisely what is quantified by information-theoretic formalism.

¹Although a general-purpose mechanism for decoding messages must be contained within the receiver, the specific instruction of which primary message to reconstruct is contained in the coded message itself. The message possesses this instruction partly due to that general capacity of the receiver. On the functional account advocated here, the same is true of any signalling system. No signal has meaning independent of sender and receiver mechanisms.

Since Bar-Hillel and Carnap [2], philosophers have been told there is a deep divide between “Shannon information” – codes – and “semantic information” – what is expressed by codes. But whether or not the source message has a semantic meaning, the encoded message certainly does, and it is a meaning that is directly relevant to the quantification of information transmitted by it. Below in section 4 I demonstrate that the first meaning of an encoded message – the information (indicative) about what the primary message was, or equivalently the instruction (imperative) how to recover it – is a kind of primitive content familiar from signalling games. It is only because the central model is a very special kind of system that its relation to other forms of communication has been neglected by philosophers. In this way, contemporary positions on the use of information theory in natural science are inordinately pessimistic.

The central model has been applied in diverse ways in the natural sciences. There has been some difficulty establishing the conditions under which it is appropriate. Below in section 6 two examples are presented, and the justification for the application of the central model in those domains is examined.

2.2 The fundamental problem

Shannon described the problem of his art as “that of reproducing at one point either exactly or approximately a message selected at another point” [46, p. 379]. A similar lesson applies here. Just because the fundamental problem does not reappear exactly in natural science does not mean the formalism of information theory cannot be applied outside mathematical communication theory. Mathematical formalism, appropriately generalised, is never so rigorously context-bound. Shannon was hopeful about the application of the mathematical theory elsewhere in natural science [45].

To see how information theoretic formalism can be generalised beyond the fundamental problem of MCT, consider the qualifier “approximately” in Shannon’s quote above. Where exactitude is not required, the system may be optimised to transmit at a rate of

‘just enough’ information. But how can we determine how much is enough unless the measurement of information has relevance for what the receiver does with it? In other words, how do I know how many bits I need unless I know what actions those bits are helping me choose between, or which states of the world those bits are helping me infer? The cost of information loss is always measured relative to the goal that information transmission subserves. Cognitive science and microbiology are applying these ideas already (see below section 6). Philosophy of information needs to catch up.

The aim of this section was to pump intuitions against the received view of information theory and its application outside MCT. The next section carries those doubts which have hopefully been raised, and soothes them by providing an inclusive understanding of information.

3 Functional interpretations of information

Functional accounts are familiar from Millikan [29], Jablonka [18], and Bergstrom and Rosvall [4, 3]. The latter work was concerned primarily with genetic information. Stegmann [50] and Rathkopf [39] have independently noted the potential breadth of the view beyond the genetic domain. The present work is a step in the same direction. Arguments that quantification of information is best approached with reference to a user can be found in Millikan [29] and Scarantino [43].

One strand of the philosophy of information aims to unpack the notion of uncertainty at the foundation of mathematical definitions. One might think ‘uncertainty’ requires the role of a user. However, Dretske, Skyrms and others place information-theoretic formalism on a foundation of objective probability. They argue that information is a user-independent quantity which, like any other natural resource, can be leveraged by organisms to help achieve their goals. In MCT, uncertainty is a quasi-mathematical, quasi-engineering concept. Mathematical formalism does not provide its own interpretation, so we have no *prima facie* reason to constrain Shannon’s definitions and theorems to their

original domain. However, we do need good reason to apply them elsewhere. Here I argue that the first major attempt to adopt information theory in philosophy, due to Dretske, was inappropriate. It focused on the wrong aspect of MCT information – its relation to probability – and ignored its relevant functional characteristics. I counsel an alternative route.

3.1 The causal interpretation

Dretske argued that both the mathematical tools and the central model that interprets them are universally applicable. Any medium through which correlations are borne is a channel. Any source of correlation – biological, artificial or inert – is a transmitter. And anything capable of interpreting that correlation – given a plausible construal of “interpretation” – is a receiver. His target was a naturalistic foundation for epistemology. He was not concerned with information per se but the tractability, in a sense crucial for methodological naturalism, of information by the scientific method. Central to Dretske’s account of belief was a distinction between “genuine cognitive systems” and mere “conduits of information” [10, p. 172]. Only the former are capable of bearing semantic content. In the 1970’s when the book was written, there was no foundational account of content (except perhaps Lewis’s *Convention* [24], and even that focused on rational actors). There was not yet a subpersonal notion of content, one that could be studied from the design stance. As argued below, in line with the theoretical work of Dennett and Millikan, low-level content is a suitable concept for natural science. Its formal underpinnings are found in evolutionary models designed by Skyrms and colleagues [49].

Dretske’s interpretation, which was to become the canonical account of causal information, made “information” extremely broad. To distil the narrower notion of semantic content he showed how extraneous properties are winnowed out in the process of belief formation. Neander [35], Skyrms [49, §3], Kraemer [20] and Scarantino [43] are among the neo-Dretskeans pursuing a similar line. From the perspective advocated here, in contrast,

information should start narrow. Functional information is already narrow: it entails only what the receiver's functional state is sensitive to, both actions and states of the world. The following subsection outlines a functional generalisation of the basic tenets of information theory, rendering it suitable for use across natural science.

3.2 A faithful generalisation of the central model

Causal interpretations, insofar as they cast information as independent of any user, face the reference class problem (see Harman's commentary on Dretske's Précis [11], and Dretske's reply). Millikan [29] and Scarantino [43] target the problem in their accounts. Scarantino points out that information can be quantified with respect to a hypothetical user. Unfortunately, he treats users as essentially epistemic. More basic than inference, contests Millikan, is function. Information is interpreted in light of what a user's functional behaviour depends upon, as opposed to what the user believes. Most natural science takes place at the level of the design stance, making use of information without reference to epistemic concepts. Relativising informational quantities to function performance allows us to retain Scarantino's solution to the reference class problem (which is itself an extension of Millikan's solution [29]), without commitment to the epistemic complexity of potential users. In line with that approach, I claim that the minimum machinery needed to specify a non-arbitrary reference class is a proper function (in the sense of goal-oriented behaviour) of a potential user. There is a measure of function performance that increases upon receipt of relevant and useful information. Received information is quantified with respect to this increase. The approach accounts for biological practice in the most flat-footed way possible. Information talk in biology describes a real, quantifiable relation.

It follows that there is a principled and useful generalisation of the central model that differs from what Dretske had in mind. He emphasised the notion of probability-raising, casting the ultimate problem of information as epistemic. While information does raise the probability of states of the world, it also raises the favourability of actions available

to its receiver. Information is a practical as well as theoretical resource. Further, the problem of reconstructing a message has no central place on this view. The fundamental problem which information helps solve is that of behavioural coordination. By discarding the causal interpretation in favour of a functional one, we are left with an account broadly in accordance with the sender-receiver framework.

A little more must be said on the relation between the fundamental problem of MCT and the more general problem of behavioural coordination. Reconstructing a message is a special case of behaviour, whether of an evolved organism or designed machine. In MCT, reconstructing a message to a desired level of accuracy is the goal of behaviour. Insofar as information transmission is a measure of how accurately a message can be reconstructed, it is more generally a measure of how accurately a piece of behaviour can be performed. Information transmission measures how much more efficiently or accurately a receiver can achieve its goals than if it received no information at all.

In MCT, lexicons are tidy sets with well-defined probability distributions. The space of all possible messages has an entropy, and information is defined derivatively on this property. Receiving a transmission reduces uncertainty about which message was selected. If we replace the epistemic notion of uncertainty about the selected message with a functional notion of accuracy, we get a generalised interpretation of the same equations. Suppose the receiver were simply a box for converting the received code back into the primary lexicon. When the received code does not uniquely determine a message, the box must choose from among the possible messages with probabilities determined by what has been received. On this interpretation, the decoding box has a space of possible actions corresponding to available primary messages. Action space is made more certain – sharper, less entropic – by receipt of a coded message. MCT applies to this case with a functional understanding of ‘uncertainty’ in place of an epistemic one.

Dretske generalised MCT in the wrong direction. He sought to retain an abstract epistemic notion of increased probability. He should have opted for a functional notion of increased accuracy.

To reiterate, reconstruction of a message is an example of a goal than can be achieved more or less accurately. There are many others. For an example in the biological realm, consider honey bees. Foragers in eusocial colonies communicate the location of food by dancing inside the hive. Hivemates who might otherwise search randomly for food are guided towards reliable food sources by following the dance. These recruited foragers are more accurate, on average, as a result of responding appropriately to their compatriots' behaviour. It is this notion of increased accuracy I contend is an appropriate generalisation of MCT.

A similar argument is made by Lean [22, pp. 239-40]. We can generalise communication to "adaptation" of the receiver to the sender, rather than identification of a reconstructed message with a source. In the examples just described, two conditions are important. The first is that action-space is measurable in terms of its favourability. It is this that is changed by the instructional aspect of information (see below section 4). Second, sender and receiver are in full cooperation, whether codesigned or coadapted by natural selection. I do not here consider the case of divergent interests, though it is crucial to the debate over animal communication. The example above used honey bees, who can be fairly idealised as cooperating fully in a foraging task. Other cases of animal signalling are not so simple. Perhaps further generalisation of information-theoretic results will yield understanding of the link between communication theory and game theory. Figure 1 describes relations between theorems of MCT. From left to right, each arrow denotes a more general theorem derived from a special one. Each is obtained by allowing parameters held fixed in the special case to vary. Speculatively, something similar will happen when we generalise further, allowing the interests of sender and receiver to vary.

I now survey three areas, two conceptual and one empirical, that recommend a generalisation of MCT centred on a functional interpretation of information.

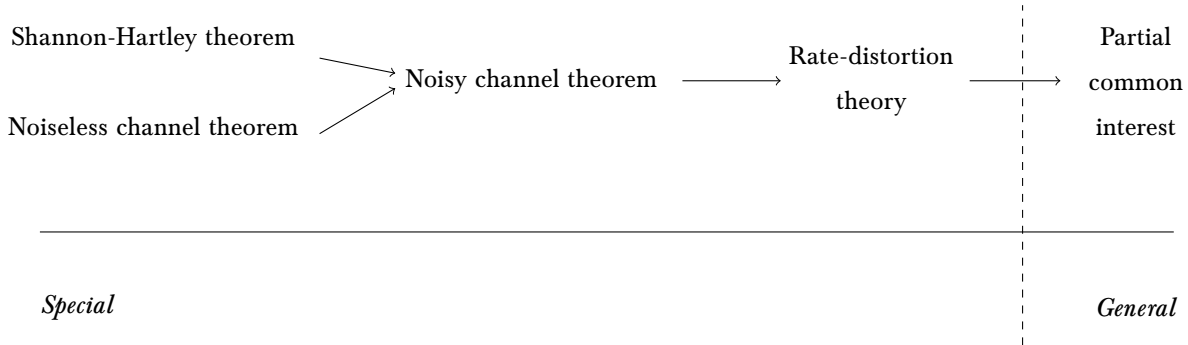


Figure 1: Relationships between formal aspects of information theory. Arrows designate the relation from special to general case. The dashed line suggests a boundary between MCT and game theory.

4 First result: information and content

Emphasis on the indicative aspect of information and a focus on its potential user’s knowing (rather than doing) mutually support each other. We can de-emphasise both simultaneously by bringing out the equally significant instructional or imperative aspect of information and its potential user’s action. In line with recent work on subpersonal content, epistemic metaphors are displaced [47, 4].

A growing (perhaps by now established) trend in naturalistic epistemology is the use of a concept of content that does not require personal-level intentional states. Its value lies in its explanatory role describing the behaviour of organisms and artificial devices too simple to be ascribed personalities [30, 47]. The classic example of this move is Skyrms’s appropriating of Lewis’s analysis of content in game-theoretic models [49]. Lewis [24] applied a mathematical framework to study the behaviour of rational actors, extracting a notion of content from the dynamics of behaviour observed in such games. Skyrms demonstrated that rationality on behalf of the actors is unnecessary. The same notion of content – the same explanations and descriptions of behaviour of the players – can be put to work when agents in the game are interpreted as evolutionary rather than rational. As with the theoretical approach spearheaded by Millikan and Shea, this move can be

regarded as a demonstration that concepts previously developed for the intentional level are applicable at the design level too.

Consider the following similarity between the formal concept of information and subpersonal content. One of the important aspects of content, discussed by both Millikan and the game theorists, is its dual aspect of indicative and imperative conditions. Signals – intentional signs – can say both how things are and what is to be done. When it comes to information we tend to implicitly adopt a purely indicative stance. In turn, we favour an epistemic interpretation of its use, on which information is what it is only with respect to some knower, independent of how they might act in response. They may be actual or hypothetical, but it is their actual or possible knowing that lends an indicative flavour to “information”. I contend that once we focus not just on knowing but on actual or possible *doing* we regain the instructional aspect of signals, thus moving closer to the contemporary understanding of content. Where 1 bit of information allows the receiver to infer one out of two equiprobable states of the world, 1 bit of instruction allows the receiver to choose one out of two equifavourable acts.

First an example. In terms of subpersonal content, an encoded signal in the central model of communication theory is *primitive*. Primitive content denotes a signal that is equally indicative and imperative. We can see that this is true of the encoded signal in the central model by overlaying a sender-receiver model on the inner portion of the Shannon system, taking the sender to be the encoding transmitter and the receiver to be the decoding receiver. The signal is then the stream of encoded symbols. Sender and receiver are dealing with encoded information whose ‘meaning’ is the primary message. The information can equally be seen as telling the decoder what the source message is (informative) and telling the decoder which target message to construct (instructional).

A test for primitivity is deliberation [24, p. 144] [16, p. 410]. Where receivers are permitted to deliberate over the action they will perform, the signal they receive has a more indicative flavour. In contrast, where senders are permitted to deliberate over what signal to send, it seems they are telling the receiver what to do. (If both are permitted to

deliberate, the correlational link between world and act is in danger of being destroyed, and communication breaks down.) In the present case, neither encoder nor decoder deliberates. Neither draws on information outside the channel through which the signal in question is flowing. The encoder takes the message as input and produces a stream of code. The decoder takes that code as input and produces a message. Since the same primary message cannot give rise to a different code, the encoder does not deliberate. And since the same code cannot be translated into different messages, the decoder does not deliberate. Finally, since neither deliberates, the code is an instance of primitive content. It is worth noting this holds regardless of the meaning of the primary message. Indeed, the primary “message” need not be composed of syntactic symbols. It need only be an element or sequence selected from a probabilistic distribution. Similarly, the target need only be an element or sequence drawn from a distribution, and it need not be the same distribution as the source. As mentioned above, Lean [22, pp. 239–40] makes a similar point, forging a path for the adoption of informational models in other domains.

Compare Piccinini and Scarantino [38, p. 19]: “Shannon’s messages need not have semantic content at all – they need not stand for anything.” (Compare also Owren et al. [36, p. 761].) Taken literally, this is false. An encoded message in the central model must stand for its primary message, otherwise there can be no definition of information rate. The authors might reply that what they meant to say is that the meaning of the *primary* message is irrelevant to the quantification of information transmission. It could be a string of meaningless symbols and transmission rate would not change. But this latter point is repeatedly conflated with the much stronger and entirely unsupported claim that “Shannon information” is irrelevant for meaning *in all domains in which information can be quantified*. It is by failing to appreciate the special nature of the central model that the claim of universal irrelevance gains traction.

Back to primitive content. It is prevalent in simple systems, which is why it was christened “primitive” by Harms [15]. Where what matters is coordinated behaviour, being told *that* another agent is performing act A is equivalent to being told *to* perform act B. It is

not just because simple agents are not literally knowers that we should shift to a functional gloss. It is just that the indicative and imperative aspects of subpersonal content – the informative and instructional aspects of the quasi-engineering concept of information – are useful concepts to apply at the design level too. The causal interpretation encourages inordinate emphasis on the epistemic. Naturalistic intentionality is better off grounded in function. Both informative (indicative, descriptive) and instructional (imperative, prescriptive) aspects are of equal significance.²

Finally, a word on Skyrms, who adopts a similar approach that differs in one crucial respect. He takes the imperative aspect of a signal to be *information about the act* to be performed. It is noteworthy that even when we recognise the dual aspect of content, it is hard to disabuse ourselves of epistemicity.

To sum up: within the informational paradigm, there is something that looks like instruction or can fruitfully be interpreted as such. Information and instruction look like indicative and imperative content. We already have a comprehensive account of content – the teleosemantic/evolutionary game theory approach whose unification has recently been argued by Artiga [1] – that endorses amalgamating indicative and imperative content for simple systems. That the bottom-up approach of signalling systems coincides with the top-down approach of teleosemantics is a sign of coherence that we should embrace. Concordance is not available on the causal interpretation because it has no notion of user action. It embodies an epistemic approach that focuses on knowing rather than doing.

So far we have talked only about information in communication channels. The next section surveys information picked up from the environment.

²A recent trend seeks to distinguish two concepts I treat as equivalent. The distinction advocated by Rescorla [40] and Lean [23] runs as follows (Rescorla cites Burge [6] as inspiration; Dan Hutto also promotes something like this). Simple signalling systems carry information in the guise of *reliable correlation* (“functional isomorphism”, “Shannon information”) – tokens that correspond to worldly states in a manner sufficient for successful behaviour. But correlational information is to be distinguished from the much richer notion of *content*, which is characterised by truth conditions. If the present reader is troubled by this purported conflation, they are encouraged to read Millikan’s response [31] which I cannot improve upon.

5 Second result: supplanting the causal/semantic distinction with the natural/intentional distinction

We have already noted confusion engendered by the term “Shannon information”. Attributing Dretske’s interpretation of MCT to Shannon encompasses many of the misconceptions recognisable in the literature on philosophy of biology and natural epistemology. Shannon had little to do with what the term now connotes. For the most part, it is Dretske’s legacy. Those who endorse a distinction between causal and semantic information invoke Shannon’s name to lend authority to the claim that MCT has no relevance for questions of meaning. Used in this way, “Shannon information” threatens to beg the question against the approach advocated here.

For one who accepts the distinction between causal and semantic information, it is possible to accept or reject the utility of the former. I endorse neither. Instead, along with Bergstrom & Rosvall [4], Rathkopf [39] and others, I reject Shannon information as a faulty concept. It conflates too much, and as such tracks nothing uniquely.

There is, however, a useful distinction to be made within biological and cognitive sciences. We will borrow Millikan’s terminology and call this the natural/intentional distinction. While section 4 emphasised intentional communication, we need to talk about natural information too. In brief, natural information is not transmitted by an entity codesigned with its receiver. It is, therefore, not an appropriate domain in which to ask questions about the efficiency of information transmission. Instead, we can ask about the efficiency of the use of the information so received.

The causal/semantic distinction [38, §§4.1-2] [13, §§2-3] has at least two sources. It is firstly a mutated form of an earlier distinction between natural and intentional meaning, which may be traced back at least to Brentano and found its clearest statement in Grice [14]. Prompted by Dretske [10, 9] the distinction took centre stage in the teleosemantic literature of the 90’s [34]. The original distinction is still hard at work in Millikan’s teleosemantics [27, §§11-12], but its mutated form is misleading. A second source is Bar-

Hillel and Carnap's clarification of "information" as it appears in MCT. They distinguished the mathematical quantity from the semantic notion which is of interest to philosophers [2]. Dretske compared Grice's approach, as well as that of Bar-Hillel and Carnap, to his own project [10, pp. 241-2, n. 1 and n. 10]. Soon after, the "still imperfectly understood" distinction was cited by Dennett [7, p. 344 col. 2] and picked up by Krebs and Dawkins [21, §§4.1-2], whence it found its way into the behavioural ecology literature and prompted ongoing scepticism about the use of information theory in the study of animal signalling [37, 36, 42].

Teleosemantic theories are theories about signs and how they can be false. The natural/intentional distinction cuts between two kinds of sign, and I assume without argument signs are not restricted to minds. Of the major players in the teleosemantic debate, Dretske, Millikan and Fodor agreed on at least one thing. The distinction made by Grice between two meanings of the term "meaning" is a respectable and useful one, and captures a difference of some import between these two kinds of sign [34, p. 116]. What is more, they agreed on the character of that distinction. Intentional signs are just those that can be false. Three rings on the bus mean that the bus is full even if the bus is not full. The vervet's chatter, when given in a situation typified by the potential presence of snakes, means that a snake is present even when snakes are absent. In contrast, clouds cannot mean rain if it does not actually rain. And smoke cannot mean fire if it was not produced by fire in the normal way. Grice, Fodor, Dretske and Millikan all agree that natural signs cannot be false, though they differ in the details of how and why intentional production entails the possibility of falsity.

One might think all a communication system needs to do is leverage natural information. Then the only difference between natural and intentional information is that the latter is being pressed into service, while the former is freely given and serves no function by itself. This is a lightweight view of intentional information, which causes problems when it encourages thinking of causal information as purely descriptive. Intentional signs are not simply reports about the world, they are also exhortations for the respondent to do

something. More importantly, the distinction entails application of different mathematical models. What matters to cognitive science and biology is optimal behaviour. The question how to optimise behaviour is different for natural information and communication. In communication the signal itself is optimised. Natural information can only support optimal receiver behaviour; the sign itself cannot be optimised. Below in section 6 I present two case studies in which optimal information transmission is investigated by applying tools from MCT.

A porous barrier separates natural and intentional signs. As is well known to behavioural ecology, many signals start life as cues. Through a process of ritualization, information that is originally a by-product of behaviour becomes codified and streamlined for purposes of efficient transmission. From a game-theoretic perspective, natural information is just communicated information where the sender's payoff matrix is unrelated to receiver behaviour. Popular models of communication are able to describe natural information by setting various parameters to zero on the sender side, rendering it less player-like and more nature-like. Skyrms [49, §1] is a case in point. There are two agential players – the familiar sender and receiver – but the information provided by Nature to the sender also has sender-receiver structure. Skyrms does not explicitly describe Nature as a “player”, but in formal models it has that character. While representing natural information as a special case of communication might sometimes be useful, Millikan [27, Part II] argues that this can blind us to the important differences between the two kinds of phenomena. The most significant difference is described above: that natural signs cannot be false [14, 34]. Of course, that also means they cannot be true, although the literature often describes their status as ‘veridical’. Truth and falsity do not apply to natural signs because they do not bear correctness conditions. They do not have correctness conditions because they do not have functions. Like all orthodoxy this view has been disputed [44]. However, it is strongly theoretically supported [28, 32, 34, 33] and recent formal work seeks to justify it in light of Skyrms's treatment [5, 47]. To repeat, the traditional view that natural signs do not have correctness conditions is supported by

the teleosemantic approach to intentionality, on which having a function is necessary for having a correctness condition.

If we think game theory is the right way to interpret both kinds of information – not just to study or come to know or understand one or both of them, but to interpret both of them in terms of mathematical games – then we might well treat natural information as a rather anaemic sort of communicated information, one whose sender’s heart isn’t in it. On the other hand if we see communication as a leveraging of a natural resource – as an attempt to imitate nature who produces information effortlessly and to no end – the situation is flipped on its head. From this perspective all information is at bottom “natural”, whether it is communicated intentionally or retrieved from an uncaring environment. Lean [23] offers such an approach. Along with Millikan I promote neither view. Intentional information is not a special case of natural information (though *true* intentional signs are or contain "root signs"; see Millikan’s diagram [29, p. 145]). Nor is natural information simply a limiting case of intentional information. As argued in section 4, there are issues with failing to respect the intentional character of communicated information. The reason truth and falsity are appropriate for intentional signs is part of the reason why they are studied by game theory and communication theory, as opposed to decision theory and statistics.

In sum, the natural/intentional distinction has real import for the application of information theory in natural science. The question of what is optimised depends on what has a function. Communicated signals bear functions, natural signs do not. The causal/semantic distinction looks implausible once the framework of MCT is appropriately generalised. Residual intuitions are shifted to the natural/intentional distinction, their original home. The next section introduces two case studies and examines their informational results.

6 Third result: rate distortion theory and applications in natural science

Now the payoff. What turns on our adoption of a functional perspective? Two case studies bring out a wider lesson recommending function as a foundation for the application of information-theoretic concepts in cognitive science and biology.

6.1 Microbiology

Many single-celled organisms can sense chemical changes in their surroundings. By navigating along gradients of changing density they can find food or avoid toxins. One species, *Dictyostelium discoideum*, uses this process of chemotaxis to coordinate mass response to a lack of nutrients. When food is scarce it is beneficial to pool resources by aggregating. *D. discoideum* cells seek each other out by alternately releasing waves of chemicals and moving in the direction of greatest concentration. When enough cells aggregate, a fruiting body forms, which helps propagate spores into a more favourable environment. These become the next generation of cells, once again adopting an individual lifestyle.

During the aggregation phase, individual cells face an informational problem. They need to know the best direction in which to travel, and they need to be sensitive to external changes in order to do this. The metabolic cost of sensitivity to fine changes in gradient impedes perfect behaviour. Like many other living things, *D. discoideum* must strike a balance. It must optimise its behaviour relative to informational constraints and the requirements of behavioural accuracy. Fortunately, rate distortion theory is designed to analyse such trade-offs. At the heart of the theory is a cost function describing the penalty for misinterpreting a signal. Depending on the cost an agent is willing to incur, it is permitted to obtain smaller amounts of information. Because each situation – across engineering, computer science, cognitive science and biology – is likely to involve different trade-offs between information and cost, the function describing optimal behaviour must be calculated anew each time.

In recognition of the problem facing *D. discoideum*, Iglesias [17] describes how rate distortion theory could be applied to explain and predict optimal behaviour. He uses a nonstandard application of the central model to describe the cell's predicament. The signal is not the external chemical gradient. Instead, it is the hypothetical decision located between the cell's receptors (whose state corresponds more or less exactly to the immediate gradient) and the cell's behaviour. The central model is applied entirely within a single cell. Mathematically the approach is acceptable, since a channel need be nothing more than a probabilistic connection between two pieces of behaviour. Given this way of modelling the situation, an appropriate optimisation function would describe the amount of information required by a migrating cell to successfully reach its target, which is equivalent to how well-correlated the cell's decision should be with its external receptors. How can we estimate information rate and cost schedule in order to derive such a function? Iglesias takes a function that measures angular deviation from the direction the cell is supposed to be moving towards. Cost increases as the angle between the cell's movement and the true direction of aggregation increases.

Iglesias provides a plausible starting point for the application of information theory to optimisation problems of this kind. However, some of the details are as yet unjustified. It is not clear how sharply cost should grow as angular deviation increases. Nor is it clear how this function might change as the cell approaches its target. Presumably the cell needs to be more sensitive to information the further it is from the goal. As it gets closer, its internal bias could override transient changes of chemical gradient that would otherwise send it in the wrong direction. Iglesias considers the possibility of internal bias, but only to demonstrate how it affects the mathematics in an idealised case. Empirical work is required to determine how much and what type of bias develops during chemotaxis. This entails a methodological problem: both bias and cost function must be derived from behaviour, but both of them are unknown or only broadly guessable at the outset. Perhaps parametric models describing both functions at once can be employed to generate testable hypotheses. These are typical methodological issues faced when fitting a model to reality.

The application of information theory in microbiology is in its early days, but signs are positive it can provide a real contribution.

The next subsection describes the application of rate distortion theory in a rather different domain, namely human perception.

6.2 Perception

We just saw an application of rate distortion theory for which the output of information transmission was a piece of behaviour, namely directional orientation and movement towards a goal. Sims [48] applies rate distortion theory in a domain closer to its original home, in which the output of information transmission is another piece of information.

Sims discusses two experiments. The first examines how accurately human subjects can categorise straight lines by length (“absolute identification”). The second determines how accurately subjects can choose the longest of two lines seen one after the other (“perceptual working memory”). Accuracy in each of these tasks is impeded by the information capacity of perceptual processing. Sims applies rate distortion theory to estimate the optimum rate given the cost of inaccuracy. He places emphasis on the lack of a known cost function, and the methods by which we can estimate one. As a result of applying the theory to both experimental procedures, Sims purportedly derives new insights into human perceptual performance [48, p. 185 col. 2 and p. 190 col. 2].

Consider the absolute identification task (taken from Rouder et al. [41]). A set of lines of N different lengths were presented randomly. Subjects were asked to choose the category, from 1 to N , in which the line belonged. After each choice the subject was informed whether or not they were correct. The mapping of the central model onto perception is more intuitive than the microbiological case. The source is the perceptual stimuli – lines of differing lengths – and the output is the subject’s response. Encoding and decoding take place within the subject’s perceptual system, and the channel capacity is inferred from the probabilities of correct responses as the number of possible lines increases. Subjects seem

to reach a point at which their performance cannot improve indefinitely [48, pp. 185–6], implying capacity is limited and rate distortion theory can be fruitfully applied.

What are the cost functions constraining human perceptual performance? Sims admits they are largely unknown. But he offers resources for estimating them by attending to performance data. His figure 5 [48, p. 186] depicts three different models of increasing fit with the data, corresponding to three different cost functions of increasing complexity. The final and most accurate cost function accords with existing ideas about perceptual “anchors” used by subjects to generate best guesses. By presenting a sequence of potential models, Sims advocates something like the following methodology. We can use rate distortion theory to derive increasingly accurate estimations of the cost function guiding perceptual tasks, while at the same time providing hypotheses as to *why* those cost functions should be at work rather than some other. One of the problems with this approach is that the cost function is not the only unknown. In the second study, performance varies with the subject’s implicit estimation of source statistics [48, p. 191 col. 2]. The experimenter must use performance data to infer both the subjective cost function and the subjective statistics. The situation is similar to that for microbiology: we have two unknown parameters of our model, and only one set of data for inferring their values. As above, the situation is not insurmountable. Sims details methods for inferring appropriate models by using empirical data together with reasonable hypotheses.

Overall, Sims sees rate-distortion theory as a tool to investigate the information-processing capabilities of biological systems [48, p. 193 col. 1]. Though he deals with what are essentially informational outputs – the responses of test subjects – he emphasises the generality of goals subserved by information processes [48, p. 193 col. 1]: “The objective for biological information processing is not (merely) the communication of information, but rather the minimization of relevant costs. Information is simply a means to an end.” This approach accords with the interpretation of Iglesias presented above.

6.3 A functional account of information provides a clear interpretation of cost

We have seen two applications of rate distortion theory in two rather different domains. Iglesias and Sims analyse biological signals with unavoidable reference to their meaning. This practice is more easily explicable on the functional view than the causal approach. One way to interpret cost is in terms of biological function. The notion that some behaviours bear a cost while others are rewarded is part of what distinguishes function and disposition. While the causal interpretation has been criticised as too broad and failing to prioritise useful information over idle correlations, the functional interpretation explicitly considers the cost of inaccuracy that is a consequence of reduced information rate. Here, information rate gets its significance from the magnitude of benefit it induces. Functional information is useful *ex hypothesi*, but there are principled bounds on its utility, as with any other resource.

In communications engineering, rate-distortion curves can be interpreted one of two ways. Suppose you know the maximum cost of inaccuracy you are willing to incur. Then the curve tells you the minimum information rate you need to transmit at in order not to exceed that cost. Alternatively, suppose you know the maximum information rate you are able to transmit at. Then the curve tells you the minimum cost you can hope to incur. In biology, however, it seems cost will always come first. Obtaining information is a strategy for reaching a goal. The metabolic resources invested in gathering information depend on how much you need, which is determined by a cost schedule covering the many ways of failing to achieve the goal. The cost of failure is then traded off against metabolic cost.

Increasing information rate in a communication system plausibly imposes metabolic costs. The situation is a special case of behavioural optimisation. Here a rate-distortion curve is the correct model to describe the relationship between improvement and metabolic cost, because the means of improvement is information transmission. Mathematical tools used to describe this kind of optimisation are taken from information theory as developed by Shannon and others. One significant change is that instead of state space being

a message in a lexicon it is a biologically relevant state of affairs. The signal contains information about that state of affairs, and quantifying that information is a crucial aspect of explaining optimal behaviour. The application of information theory in biology and cognitive science concerns optimisation with respect to the statistical properties of both source and target, which need not be (and in the biological case hardly ever are) messages in well-defined lexicons.

One final point about biological cost is in order. Ultimately, the costs that shape the functions whose performance is studied in the manner described above are provided by natural selection. Environmental pressures determine optimal behaviour. Functions like microbe motility and perceptual categorisation are not performed for their own sake. They contribute to the survival and reproduction of whichever entity or entities are under selection. We should therefore expect the cost function of an individual piece of behaviour to derive from cost functions that govern selective pressures. Donaldson-Matasci et al. [8] describe evolutionary fitness in informational terms. In an analysis deriving ultimately from Kelly's interpretation of information rate [19], the fitness penalty of failing to heed information is bounded by the quantity of that information. Frank [12] offers an informational interpretation of evolutionary fitness that seems to accord with this view. For simple models at least, our concepts of information and fitness cost are deeply entangled. When fitness is interpreted as growth rate, it can be measured with a unit that is commensurable with information units [8, p. 228 col. 2] (of which the bit is the most familiar example). And when cost functions are interpreted as fitness penalties, their proper unit of measure is the same as that of information. Rate-distortion curves can then be interpreted directly as the contribution to fitness afforded by information transmission. To apply these ideas more concretely, such as to the work of Iglesias, would require an understanding of how individual behaviours contribute quantitatively to the fitness of organisms.

In sum, a functional interpretation of information accords well with a growing trend in natural science to consider the optimisation of informational processes. In contrast to the causal interpretation, biological information cannot be pulled apart from the costs

incurred to handle it and the benefits attained by using it. Information allows a principled redistribution of physical resources, entailing optimised behaviour that contributes to downstream functions and, eventually, evolutionary fitness.

7 Conclusion

The philosophy of information is heavily indebted to Dretske. It has tended to retain his mistakes along with his triumphs. Though information can be generalised in a philosophically interesting way, we should not jump straight from communication theory to epistemology. The concept of function is a natural bridge that retains the mathematical structure of information theory, and supports its use in domains such as behavioural ecology, microbiology, and cognitive science.

References

- [1] Marc Artiga. “Teleosemantic Modeling of Cognitive Representations”. *Biology & Philosophy* 31.4 (2016), pp. 483–505. doi: 10.1007/s10539-016-9525-3.
- [2] Yehoshua Bar-Hillel and Rudolf Carnap. “Semantic Information”. *The British Journal for the Philosophy of Science* 4.14 (1953), pp. 147–157.
- [3] Carl T. Bergstrom and Martin Rosvall. “Response to Commentaries on “The Transmission Sense of Information””. *Biology & Philosophy* 26.2 (2011), pp. 195–200. doi: 10.1007/s10539-011-9257-3.
- [4] Carl T. Bergstrom and Martin Rosvall. “The Transmission Sense of Information”. *Biology & Philosophy* 26.2 (2011), pp. 159–176. doi: 10.1007/s10539-009-9180-z.
- [5] Jonathan Birch. “Propositional Content in Signalling Systems”. *Philosophical Studies* 171.3 (2014), pp. 493–512. doi: 10.1007/s11098-014-0280-5.
- [6] Tyler Burge. *Origins of Objectivity*. Oxford: OUP, 2010.
- [7] Daniel C. Dennett. “Intentional Systems in Cognitive Ethology: The “Panglossian Paradigm” Defended”. *Behavioral and Brain Sciences* 6.3 (1983), pp. 343–355. doi: 10.1017/S0140525X00016393.
- [8] Matina C. Donaldson-Matasci, Carl T. Bergstrom, and Michael Lachmann. “The Fitness Value of Information”. *Oikos* 119.2 (2010), pp. 219–230. doi: 10.1111/j.1600-0706.2009.17781.x.
- [9] Fred Dretske. *Explaining Behavior: Reasons in a World of Causes*. Cambridge, Mass.: MIT Press, 1988.
- [10] Fred Dretske. *Knowledge and the Flow of Information*. Cambridge, Mass.: MIT Press, 1981.
- [11] Fred Dretske. “Précis of Knowledge and the Flow of Information”. *Behavioral and Brain Sciences* 6.01 (1983), pp. 55–63. doi: 10.1017/S0140525X00014631.
- [12] Steven A. Frank. “Natural Selection. V. How to Read the Fundamental Equations of Evolutionary Change in Terms of Information Theory”. *Journal of Evolutionary Biology* 25.12 (2012), pp. 2377–2396. doi: 10.1111/jeb.12010.
- [13] Peter Godfrey-Smith and Kim Sterelny. “Biological Information”. *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2016. 2016.
- [14] H. P. Grice. “Meaning”. *The Philosophical Review* 66.3 (1957), pp. 377–388. doi: 10.2307/2182440.
- [15] William F. Harms. “Primitive Content, Translation, and the Emergence of Meaning in Animal Communication”. *Evolution of Communication Systems: A Comparative Approach*. Ed. by D. Kimbrough Oller and Ulrike Griebel. 2004, pp. 31–48.

- [16] Simon M. Huttegger. “Evolutionary Explanations of Indicatives and Imperatives”. *Erkenntnis* 66.3 (2007), pp. 409–436. doi: 10.1007/s10670-006-9022-1.
- [17] P. A. Iglesias. “The Use of Rate Distortion Theory to Evaluate Biological Signaling Pathways”. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications* 2.1 (2016), pp. 31–39. doi: 10.1109/TMBMC.2016.2623600.
- [18] Eva Jablonka. “Information: Its Interpretation, Its Inheritance, and Its Sharing”. *Philosophy of Science* 69.4 (2002), pp. 578–605. doi: 10.1086/344621.
- [19] J. L. Kelly. “A New Interpretation of Information Rate”. *Bell System Technical Journal* 35.4 (1956), pp. 917–926. doi: 10.1002/j.1538-7305.1956.tb03809.x.
- [20] Daniel M. Kraemer. “Natural Probabilistic Information”. *Synthese* 192.9 (2015), pp. 2901–2919. doi: 10.1007/s11229-015-0692-6.
- [21] John R. Krebs and Richard Dawkins. “Animal Signals: Mind-Reading and Manipulation”. *Behavioural Ecology: An Evolutionary Approach*. Ed. by J. R. Krebs and N. B. Davies. 2nd ed. Oxford: Blackwell Scientific, 1984, pp. 380–402.
- [22] Oliver Lean. “Biological Information (PhD Thesis)”. PhD thesis. 2016.
- [23] Oliver M. Lean. “Getting the Most out of Shannon Information”. *Biology & Philosophy* 29.3 (2014), pp. 395–413. doi: 10.1007/s10539-013-9410-2.
- [24] David Lewis. *Convention: A Philosophical Study*. Oxford: Blackwell, 2002.
- [25] Olimpia Lombardi, Federico Holik, and Leonardo Vanni. “What Is Shannon Information?” *Synthese* 193.7 (2015), pp. 1983–2012. doi: 10.1007/s11229-015-0824-z.
- [26] James Maclaurin. “Commentary on “The Transmission Sense of Information” by Carl T. Bergstrom and Martin Rosvall”. *Biology & Philosophy* 26.2 (2011), pp. 191–194. doi: 10.1007/s10539-010-9233-3.
- [27] Ruth G. Millikan. *Beyond Concepts: Unicepts, Language, and Natural Information*. Oxford, New York: Oxford University Press, 2017.
- [28] Ruth G. Millikan. *Language, Thought, and Other Biological Categories*. MIT Press, 1984.
- [29] Ruth G. Millikan. “Natural Information, Intentional Signs and Animal Communication”. *Animal Communication Theory*. Ed. by Ulrich E. Stegmann. Cambridge University Press, 2013, pp. 133–146.
- [30] Ruth G. Millikan. “Pushmi-Pullyu Representations”. *Philosophical Perspectives* 9 (1995), pp. 185–200. doi: 10.2307/2214217.

- [31] Ruth G. Millikan. “Reply to Rescorla”. *Millikan and Her Critics*. Ed. by Dan Ryder, Justine Kingsbury, and Kenneth Williford. John Wiley & Sons, 2013, pp. 103–106.
- [32] Ruth G. Millikan. “Truth Rules, Hoverflies, and the Kripke-Wittgenstein Paradox”. *The Philosophical Review* 99.3 (1990), pp. 323–353. doi: 10.2307/2185347.
- [33] Ruth G. Millikan. *Varieties of Meaning: The 2002 Jean Nicod Lectures*. MIT Press, 2004.
- [34] Ruth G. Millikan. “What Has Natural Information to Do with Intentional Representation?” *Royal Institute of Philosophy Supplements* 49 (2001), pp. 105–125. doi: 10.1017/S135824610000713X.
- [35] Karen Neander. “Toward an Informational Teleosemantics”. *Millikan and Her Critics*. Ed. by Dan Ryder, Justine Kingsbury, and Kenneth Williford. John Wiley & Sons, 2013, pp. 21–40.
- [36] Michael J. Owren, Drew Rendall, and Michael J. Ryan. “Redefining Animal Signaling: Influence versus Information in Communication”. *Biology & Philosophy* 25.5 (2010), pp. 755–780. doi: 10.1007/s10539-010-9224-4.
- [37] Jessica Pfeifer. “The Use of Information Theory in Biology: Lessons from Social Insects”. *Biological Theory* 1.3 (2006), pp. 317–330.
- [38] Gualtiero Piccinini and Andrea Scarantino. “Information Processing, Computation, and Cognition”. *Journal of Biological Physics* 37.1 (2011), pp. 1–38. doi: 10.1007/s10867-010-9195-3.
- [39] Charles Rathkopf. “Neural Information and the Problem of Objectivity”. *Biology & Philosophy* (2017), pp. 1–16. doi: 10.1007/s10539-017-9561-7.
- [40] Michael Rescorla. “Millikan on Honeybee Navigation and Communication”. *Millikan and Her Critics*. Ed. by Dan Ryder, Justine Kingsbury, and Kenneth Williford. John Wiley & Sons, 2013, pp. 87–102.
- [41] Jeffrey N. Rouder et al. “Learning in a Unidimensional Absolute Identification Task”. *Psychonomic Bulletin & Review* 11.5 (2004), pp. 938–944. doi: 10.3758/BF03196725.
- [42] Sahotra Sarkar. “Information in Animal Communication: When and Why Does It Matter?” *Animal Communication Theory*. Cambridge University Press, 2013.
- [43] Andrea Scarantino. “Information as a Probabilistic Difference Maker”. *Australasian Journal of Philosophy* 93.3 (2015), pp. 419–443. doi: 10.1080/00048402.2014.993665.
- [44] Andrea Scarantino and Gualtiero Piccinini. “Information Without Truth”. *Metaphilosophy* 41.3 (2010), pp. 313–330. doi: 10.1111/j.1467-9973.2010.01632.x.
- [45] C. Shannon. “The Bandwagon”. *IRE Transactions on Information Theory* 2.1 (1956), pp. 3–3. doi: 10.1109/TIT.1956.1056774.

- [46] C. E. Shannon. “A Mathematical Theory of Communication (Part 1)”. *Bell System Technical Journal* 27.3 (1948), pp. 379–423. doi: 10.1002/j.1538-7305.1948.tb01338.x.
- [47] Nicholas Shea, Peter Godfrey-Smith, and Rosa Cao. “Content in Simple Signalling Systems”. *The British Journal for the Philosophy of Science* (2017). doi: 10.1093/bjps/axw036.
- [48] Chris R. Sims. “Rate–distortion Theory and Human Perception”. *Cognition* 152 (2016), pp. 181–198. doi: 10.1016/j.cognition.2016.03.020.
- [49] Brian Skyrms. *Signals: Evolution, Learning, and Information*. Oxford University Press, 2010.
- [50] Ulrich E. Stegmann. “On the ‘Transmission Sense of Information’”. *Biology & Philosophy* 28.1 (2013), pp. 141–144. doi: 10.1007/s10539-012-9310-x.