Top-down causation by information control: from a philosophical problem to a scientific research programme

G. Auletta¹, G. F. R. Ellis^{2,*} and L. Jaeger³

¹*Pontifical Gregorian University, 00187 Rome, Italy*

²*Mathematics Department, University of Cape Town, Rondebosch 7701,
Republic of South Africa*

³*University of California, Santa Barbara, CA 93106, USA*

It has been claimed that different types of causes must be considered in biological systems, including top-down as well as same-level and bottom-up causation, thus enabling the top levels to be causally efficacious in their own right. To clarify this issue, the important distinctions between information and signs are introduced here and the concepts of information control and functional equivalence classes in those systems are rigorously defined and used to characterize when top-down causation by feedback control happens, in a way that is testable. The causally significant elements we consider are equivalence classes of lower level processes, realized in biological systems through different operations having the same outcome within the context of information control and networks.

Keywords: functional equivalence class; information control; information selection; operation; sign; top-down causation

1. INTRODUCTION

At the most general level, the issue of top-down causation can be seen as the problem of how higher levels of reality (roughly, levels of greater complexity¹) can possibly have any causal effectiveness on lower levels (Simon 1962; Campbell 1974). If it is assumed (according to the standard view in physics) that purely physical effects determine what happens at the lower levels and thereby also completely determine what happens at the higher levels, how can there be any real meaning to higher level causes and effects?

Many scientists consider ‘top-down causation’ to be unreal—they believe it is just a complicated way of describing things which in the end confuses the real causal patterns, which are believed to be bottom-up only (see figure 1*a*). It is also assumed that phenomena that are not easily understandable in a bottom-up way today will be well understood in the future. This approach has been extended to all natural systems thanks to the huge success of the application of reductionist methodology in physics and, in recent decades, in molecular biology and neuroscience. According to Francis Crick’s famous dictum: ‘You, your joys and your sorrows, your memories and your ambitions, your sense of personal

identity and free will, are in fact no more than the behaviour of a vast assembly of nerve cells and their associated molecules’ (Crick 1994). The emphasis in the phrase ‘no more than’ is a denial of the reality of anything additional to the pure assembly of cells and is therefore also a rejection of top-down causation.

A similar point of view comes from some emergence theorists, who suggest that since complex systems can self-assemble from the bottom up, in biological systems there is no need for influence of the whole on the parts (Holland 1997). In both cases, the suggestion is that the introduction of top-down causation is misleading; the real causal powers reside in the lower level physics and associated classical chemistry; they alone determine what happens.

Nonetheless, there is a wide literature on the emergence of autonomous higher levels of complexity and the role of top-down causation in the hierarchy of complexity (see Clayton & Davies (2006) and Murphy & Stoeger (2007) and references therein), expressing a need felt by many scholars to overcome traditional reductionism. Unfortunately, this discussion is mostly confined to philosophical considerations that have changed little the way scientists (especially physicists and molecular biologists) consider their own work.

In this paper, we try to refine relevant concepts in order to translate a philosophical examination of top-down causation into a scientific programme able to make predictions and experimental tests. As we shall see, there are already experiments in this direction, though they have not been interpreted in the way we

*Author for correspondence (ellis@maths.uct.ac.za).

L.J. wishes to dedicate this paper to the memory of Father Benvenute Bavaro, OFM.

¹Leaving aside its formal definition as that goes further than the scope of this paper.

Figure 1. Bottom-up, top-down and same-level modes of action. (a) The classical standard view understands everything in the universe as happening in terms of bottom-up action only, so that efficient causation is seen essentially as bottom-up. (b) A more careful view is to consider the bottom-up mode of action as providing a space of possibilities, the same-level mode (dark grey arrows) as the true dynamic causation, and the top-down mode (light grey arrows) as changing the causal relations below. Dynamic causation (same-level) can be both efficient and circular. Here, the circular one is shown, since it plays a more crucial role in the context discussed here. While not represented in (a), circular causation can also be part of 'bottom-up action only' models.

articulate here. The inquiry regarding top-down causation is somewhat different if we consider cases where consciousness and intelligence are involved, as distinct from those where they are not included. Additionally, important distinctions occur between cases where life is involved, and where only physical and chemical interactions occur. This paper will focus on the case of life at its most elementary level, but not on issues raised by intelligence and intention. We think that the concepts presented herein could be crucial for systematically addressing the problem of emergence of complexity and related evolutionary aspects, which we will consider elsewhere.

2. A RECENT EXPERIMENT

A European team (Wegscheid *et al.* 2006) performed *in vivo* experiments of complementation on the bacteria *Escherichia coli* and *Bacillus subtilis*. These bacteria have distinct RNase P enzymes: type A for *E. coli* and type B for *B. subtilis*, respectively. Interestingly, these RNases P have significantly different three-dimensional architectures (figure 8) that are associated with important biophysical and biochemical differences *in vitro*. For instance: (i) *B. subtilis* P RNA 5' and 3' ends are autolytically processed after association with its protein subunit, while this operation is enzyme dependent in *E. coli*; (ii) RNase P from *B. subtilis* but not *E. coli* forms dimers consisting of two RNA and two protein subunits; (iii) *B. subtilis* RNase P holoenzyme binds to pre-tRNA with a much higher affinity than mature tRNA, whereas this difference in affinity is attenuated in *E. coli*; and (iv) type A and type B RNase P RNAs have distinct metal ion requirements.

Now, despite these structural and functional differences *in vitro*, it was shown that types A and B can replace each other *in vivo* without loss in functionality, at least

under standard growth conditions. It is interesting to stress that the lower physical-chemical stability of hybrid holoenzyme complexes does not raise functionality problems. Already this shows that the traditional reductionist point of view is insufficient. Indeed, the concept of function is now widely acknowledged in biology (see also Collins *et al.* 2007).

The authors of the experiment interpret this situation by making the hypothesis that there could be 'conserved features' of bacterial RNase P RNA and protein subunits essential for their primary functional activity *in vivo*. While this fact is difficult to deny based on the reality of a consensual core structure for RNase P enzymes in all living organisms (Altman 2007; Kirsebom 2007), we, however, think that this explanation is only partial and hides a greater truth. In our opinion, their results can be seen in a different light, by considering top-down effects through the information control that the unicellular organism exercises on and through these functions. In the following, we propose this new conceptual framework.

3. TOP-DOWN CAUSATION BY INFORMATION CONTROL

Generically, top-down causation works by higher level conditions setting the context for lower level processes (see Campbell 1974; Bishop & Atmanspacher 2006; Clayton & Davies 2006; Deacon 2006; Ellis 2006a,b; Murphy & Stoeger 2007). Already in physics and chemistry there are boundary conditions that act on lower level processes. For this reason, top-down causation *applies whatever level is chosen as the reference level*, for there is no known ontologically preferred level (as we do not even know what the lowest level is, there is no known fundamental level). However, in the cases of physics and biology mentioned above, it is not always clear whether traditional reductionist point of view is actually overcome, since these boundary conditions can again be understood as a complicated effect of much more elementary processes.

Therefore, in this paper, we focus on a specific and stronger kind of top-down causation that plays a very relevant role in biology: top-down causation by *information control*. This does not mean that we think this is the only way top-down causation is probably relevant within the biological domain (some other possibilities are discussed in Ellis 2008). We shall show that top-down causation by information control may be clearly found in biological organisms already at the most elementary biomolecular level, suggesting that there is no possibility to reinterpret these results in reductionist terms.

3.1. Some basic definitions

Top-down causation by information control is the way a higher level instance exercises control of lower level causal interactions through feedback control loops, making use of functional equivalence classes of operations. At the biological level, this consists of applying controlling signals from the higher level to influence the proper mode of action of lower level

items (figure 1b). Here, a signal is any variation or pattern in a physical or chemical medium that can convey information or be treated as a sign. If we succeed in showing meaningful evidence that this stronger form of top-down causation already happens at the bio-molecular level, we have also succeeded in showing that this occurs in all living systems. Even stronger versions of top-down causation can be found when intentionality and free will are involved; however, as already pointed out, we confine our examination here to the most basic biological domain.

A key concept here is that of *equivalence class of lower level operations*,² discussed in §4, where operations occurring in biological systems can be considered as coordinated space-time pathways of physical–chemical interactions. The criterion for an equivalence class of operations is the nature of the *outcome* of an operation relative to an established goal—if two different operations give the same outcome, they can be considered equivalent. Thus the concern here is about *functional* equivalence classes (sets of operations that produce the same outcome).

Another central point is, of course, information control (§5). In general, information control is the ability to use signals to attain or maintain a specific goal. When different forms of information control are possible, functional equivalence classes are what really counts, since they gather together the possible different operations by which the goal can be attained. Let us stress that any control mechanism is ultimately control in terms of information, even though it makes use of some energy or material to convey that information. Information is consequently a dominant factor of life (Küppers 1990; Rasmussen et al. 2004; Roederer 2005). Our point is that the demonstration of the constitution and keeping of functional equivalence classes through information control—coming from an above instance—can be seen as strong evidence of top-down causation in biological systems affecting lower level modes of operations from a higher level of functional organization.

3.2. On the nature of causes

The general nature of causation is a key issue in considering these topics. The majority of scientists think that top-down causation means to act directly on a lower ontological or less complex level of reality without the intermediary of causes acting at this same level. Indeed, a majority of both top-down-causation supporters and detractors assume that it substitutes for the specific modes of action at lower ontological levels. It is easy to see that in this way, the principle of the closure of the physical world from the point of view of action and interaction would be violated in most cases, at least in those where actions according to goals are involved.

To achieve this aim, considering the nature of causes and especially distinguishing between *dynamic* and *non-dynamic causes* (Auletta 2008c) is of primary

importance. We may say that all causes have causal power, but only dynamic causes have causal effectiveness. *Effectiveness* means that the causal agent, in ideal condition, can positively give rise to a certain effect through interactions at the same ontological level, as when a ball moving with a certain speed is able, after collision, to set another ball in motion. *Power*³ here means having a role—although *not* an *effective* one—in the production of something. In other words, causes provided with causal power, but without effectiveness, influence the realization of a certain result. But, even in ideal condition, they are not sufficient for *producing* that result since at least an effective factor must also be at work to actually achieve some production. We have effective causal interactions when some dynamic physical magnitudes (like energy or momentum) are involved. This can be understood, in general, through the so called ‘transference theory of causation’ (Fair 1979; Salmon 1984; Dowe 1992, 1995; Salmon 1994).

In our language, in fact, dynamic causes are same-level *efficient causes*, like thermal energy determining the melting of ice, or *circular causes*, which are those present in nonlinear, self-increasing phenomena, like those commonly occurring in autocatalytic chemical reactions, for instance when a chemical X is involved in its own production as in

where k_1 and k_{-1} are reaction rate constants.

Non-dynamic causes, on the contrary, are essentially of two types: *material causes* (from below) that represent the support for processes and entities at a higher level of complexity and can therefore be considered as possibility conditions for those processes and entities (for instance, the various chemicals constituting the biomolecules and entering chemical interactions underlying biological phenomena); and *formal causes* (from above) that are restrictions of the space of possibilities (for instance, during DNA replication, only certain chemical reactions are able to occur owing to the context).

Top-down causes are causes that *do not act* at the same ontological level (figure 1). They, notwithstanding, at least when information control is exerted, concur to produce certain effects when they are combined with dynamic causes at a lower level. Since top-down causes, in this context, also involve dynamic causation at the lower level, they *do* have causal effectiveness. However, since they do this *through* the specific mode of dynamic causation of the lower levels, and since dynamical causes are here understood as respecting physical conservation principles (this holds true both for efficient and circular causation), there is no violation of the closure of physical causation. Indeed, causal closure, from our point of view, is not broken if and only if you shall not introduce from the exterior some causal agency that modifies some physical interactions in a

²This notion was already introduced into psychology by Lashley (1942) and into neural sciences by Hebb (1949, pp. 38–59), and see also Pearl (1998, 2000).

³Causal power has obviously no relation with the dynamic physical magnitude defined as work per unit of time.

Figure 2. The hierarchical organization of the structure of an RNA biopolymer. From linear sequence information to two-dimensional and three-dimensional networks based on specific intra-molecular recognition. RNA biopolymers are informational and functional molecules that have the ability to fold and assemble into highly intricate three-dimensional architectures that take advantage of networks of specific tertiary interactions (Lescoute & Westhof 2006).

way that represents a violation of physical conservation laws. To avoid such a violation it is sufficient to deny that the non-dynamic causes are responsible for the exchange of an additional amount of some conserved quantity.⁴ In conclusion, a careful application of the principle of closure of the physical world to different levels of complexity should lead to the result that it is impossible to *dynamically* act from one ontological level to another, either from above or below; nevertheless, action according to goals actually comes into play already at the most basic level considered here (that of molecular biology).

Top-down causes can therefore be considered as a combination of formal causes from above, material causes from below, and operations embedded in circular causes (feedback circuits) at the middle ontological level (see figure 1b). In a biological context, they can very often be understood as a *teleological* causation, since goals play a decisive role as far as information control is involved. In fact, the equivalence classes coming into play in information control are precisely characterized by the goals to be attained (see §4.2).

4. FUNCTIONAL EQUIVALENCE CLASSES

4.1. Operations and their conditions

Operations are the elements of the functional equivalence classes considered by us, here biochemical operations. Yet, an operation cannot be reduced to the pure (low-level) physical-chemical interactions *per se* but is rather a space-time pathway of such interactions;

different pathways therefore define different operations. (Interactions are physical-chemical exchanges of physical magnitudes like charge, mass, energy, etc.)

There are three conditions for having such operations in biology.

- (i) *A space of alternative possibilities* from below (the material causes), without which one cannot have a set of different functionally equivalent operations. This set of multiple possibilities exists for microbiological processes, since biopolymers are sufficiently complex to enter into—and be integrated in—very different forms of biological processes. For instance, catalysis can be performed by either RNA or protein molecules, two very different types of biopolymers.
- (ii) *Information selection*. In order to have a specific operation, we have to select some elements from the possibility space. This is actually a selection of information, since the elements involved are biomolecules whose primary structure indeed encodes information (Lehn 2004). Information selection is a complex modality of information that is already present at the physical level, where it occurs at the most elementary level when there is interplay between at least two quantum mechanical systems open to an environment (Auletta 2005, 2006). The elements indispensable to information selection are the following. (a) *Mutual information*: when several systems share a common pool of information, for instance, when there is a coupling of the receiver with the input source. Formally, the mutual information between system 'A' and system 'B' is defined as the entropy of A plus the entropy of B minus their joint entropy, or $I(A : B) = -(H(A, B) - H(A) - H(B))$. (b) *A source of variety in the input information*. This

⁴Note that, in general, *non-violation* of something does not imply *reduction* to it. Thus, the fact that non-dynamic causes do not violate the physical conservation principles (and therefore the causal closure) does not mean at all that they should be reduced to purely physical (effective) interactions, which are the embodiment of those principles.

Figure 3. Equivalence classes, operations and functions. (a) Two functionally equivalent networks. Different nodes and contacts specify different operations. Modularity is readily apparent as some local nodes and related contacts are found at multiple locations within each network. The whole network can be seen as a hierarchical organization of smaller interchangeable lower order networks. Open circles, nodes; black lines, two-dimensional contacts; red lines, three-dimensional contacts. (b) Examples of operations and functions (see also §2). If the two networks in (a) are taken as the three-dimensional structure of two RNA biopolymers that belong to the same functional class (e.g. RNase P RNA), then some of their constitutive structural modules belong to lower level functional equivalence classes.

source may be very ordered (as is the case for quantum systems showing zero entropy) or more disordered. Moreover, such a source of variety is very often due to random events, like many mutations in DNA, or even during epigeny or ageing, and therefore plays an important role in biology. Even in cases where such variety is not random, what is important is that it is out of the control of the system it enters. This variety may be due to quantum uncertainty, or just statistical fluctuations based on the fact that biological systems are made of many billions of discrete low-level units. The great American philosopher Charles S. Peirce (1931–1958, CP 1.159, 1.174, 1.302, 1.405, 5.119, 6.30–32, 6.57–59, 6.64) emphasized that in the world there is not only uniformity, but rather the principle of variety is ‘the most obtrusive character of the universe’, which ‘no mechanism can account for’. (c) A choice (decision) event selecting one of the incoming information variants. In the simplest case, this information selection event has the form $I_k = -\lg p_k$, where p_k is the probability that the k th event occurs. This might take place in a random way as in quantum measurement, where a component of the initial superposition state (the source of variety) is chosen randomly. In the biological domain, of all the environmental data (the source of

variety) only some are selected according to some criterion of relevance. In the case of microbiology, this is often through recognition of some biomolecules by receptors and rejection of others, for example, by neurotransmitter receptors in axons. This choice, however, being made between the options provided by the variety of incoming information, conveys some knowledge about the source system.

Summarizing, in all its generality we define information selection as *a decision event that selects particular input information through coupling with it*. As we shall see below, we have a second degree of information selection, and therefore true top-down causation when a controller makes a guess and selects relative to a certain goal. Here we also need a second-level information theory, semiotics (§5.3).

- (iii) *Modularity*. This means there are units that are somehow *uncoupled from the external world* in that they allow information hiding, encapsulation, hierarchy and abstraction (Booch 1994). As internal variables are hidden, external inputs do not always determine a unique reaction of a module, for their effects depend also on its internal state. Modularity, therefore, represents a sort of divergence in the effects. The internal

operations of each module are then effectively decoupled from any other same-level module. This allows that an operation can be informationally controlled *without* being dynamically causally affected by other modules or foreign factors. Modules are connected in networks that are the operative contexts where top-down causation properly takes place. It may happen that a certain network can be a sub-network (module) of a more complex network (see also figure 3a).

4.2. Networks and operations

Although the notion of network is used in a huge range of disciplines in the biological domain (see Barabasi & Oltvai 2004; Schmid & McMahon 2007), it is usefully applicable to biochemistry only from the supramolecular level upwards (Lehn 2004) and fits very well with our examination of equivalence classes and information control. Here, we employ a general notion of network understood as a cluster of interdependencies among units called nodes, without introducing specific assumptions about the nature of either the relations or the nodes.

During the constitution of networks within biological systems, information selection indeed occurs—many different molecules (the space of possibilities representing the input information) come to be coupled, becoming in this way nodes of a net in which ‘choice’ events are initially produced, giving rise to more significant nodes (hubs) that may be considered as selecting (and gathering) information from other nodes. When this network becomes modularized, it may instantiate a specific operation.

For example, biomolecular compounds such as RNA and proteins are structured entities showing a high degree of internal interconnection (see figure 2). One can see RNA assembly and folding as a process of hierarchical and stepwise formation of a final network (the tertiary structure) through information selection, requiring first formation of specific helical elements through Watson–Crick base pairing. This defines the RNA secondary structure that is formally a network of hydrogen bonds occurring between complementary nucleotide residues (adenine (A) goes with uracil (U) and cytosine (C) goes with guanine (G)). In a second step, the presence of salts in the medium triggers the collapse of the secondary structure into structural intermediates through partial neutralization of the negative ribose phosphate backbone. Lastly, tertiary motifs that specify for particular tertiary contacts lead to the final native fold that is able to perform specific operations and therefore can carry inter-molecular recognition, catalytic or mechanical functions. As shown in figure 2, the final three-dimensional fold of an RNA molecule can be seen as a complex network of non-covalent interactions occurring between distant sites (or nodes) within the linear polymer sequence (Lescoute & Westhof 2006).

The folding of biomolecules can thus be formalized as a network in which several subunits can be individuated, each of them playing a definite role within that wider context. The various subunits are determined through information selection (which picks up

elements from the space of possibilities) and, spontaneously interacting, constitute a concrete complex of interrelations. Hence, the subunits are the nodes of the network and the array of interdependences among nodes specify a determinate pathway of physical–chemical interactions, that is, a definite operation (e.g. figure 3a). Finally, thanks to modularity, such an operation (represented by the network) is shielded against external perturbations, and then controlled and reiterated, in a top-down fashion, by a higher level network.

Another example of a network is provided by inter-actome pathway maps, and as an isolated operation we may consider endocytosis, the operation through which some external molecule will be engulfed in a cell (Schmid & McMahon 2007).

In many cases, some of the nodes of a network can be substituted by different suitable sets of chemicals without losing the overall features of the network itself. It is noteworthy that some nodes can be simply dropped without altering the effect of the whole pathway. This shows that the same function can be instantiated through different clusters of physical–chemical interactions.

4.3. Equivalence class definition

Mathematically, an equivalence relation is a type of relation on a set that provides a way for elements of that set to be identified with other elements of the set. Those elements considered equivalent through this identification form an equivalence class.⁵

Let W be a set and let w , x and y be elements of W . An *equivalence relation*, $\sim$, on W is a relation on W , that is,

Reflexive. w is equivalent to w for all w in W .

Symmetric. If w is equivalent to x , then x is equivalent to w .

Transitive. If w is equivalent to x and x is equivalent to y , then w is equivalent to y .

If W is a set, w an element of W , and $\sim$ an equivalence relation on W , the *equivalence class* of w is the set of all elements of W equivalent to w under $\sim$.

4.4. Functional equivalence classes

As we have said, the concept of function is now widely acknowledged in biology. However, we emphasize that a function defines an equivalence class. When one speaks of equivalence in general, one understands a ‘possibility of substitution in any context’. It is fundamental that functional equivalence classes are on the contrary context sensitive. The concept of equivalence class, which is the object of this paper, is therefore not a formal logical construct but a pure functional biological category, where different operations are considered functionally equivalent if they produce the same outcome for some functional purpose (the goal). Thus,

⁵Equivalence classes (see http://www.iscid.org/encyclopedia/Equivalence_Class) are based on equivalence relations (see http://www.iscid.org/encyclopedia/Equivalence_Relation).

Figure 4. Splicing. Splicing as an example of equivalence classes. Here, the fact that the different modes of operations indeed provide mRNA (as an outcome, and independently of the operations that mRNA can, in turn, perform relative to a further function) points out that these modes pertain to an equivalence class.

we focus on functional equivalence classes rather than purely formal equivalence classes, even if the properties defined previously must also hold for them.

To be specific, let us consider the biochemical subsystems that are required to synthesize a minimal cell: genome replication, transcription, translation, correct protein folding and any necessary post-translational amino acid modifications (Forster & Church 2006). Some of these abstract functions can be fulfilled by different pathways of physical-chemical interactions, i.e. by different operations (see figure 3). When different operations fulfil the same function, they form an equivalence class (in this case, the formal relations of §4.3 will be fulfilled). Since operations equivalent with respect to a certain function *are not* automatically equivalent for other functions, it is important to identify unequivocally the function concerned. Therefore, the criterion by which items are judged to be or not to be members of such classes is *only* a specific function. In this way, functional equivalence classes are characterized by a *part-whole* relation, which, in turn, links a particular kind of function directly with the issue of signs (see §5.3). It should be noted that good examples of biological systems belonging to a particular functional equivalence class are evolutionarily unrelated functional biomolecules that result from convergent evolutionary processes. A clear example of functional equivalence class is provided by splicing, which is the operation involved in the transition from pre-mRNA to mature mRNA. The function 'splicing' is performed through different operations (making use of different chemicals, see figure 4). Despite such a difference, the outcome is the same as far as the splicing is concerned.

Functional equivalence classes can contain multiple modules belonging to functional equivalence classes of lower complexity (figure 3). The demonstration

that lower functional equivalence classes are under the hierarchical control of higher level ones can be seen as evidence of top-down causation. We recall that in order that a higher level network can be considered as constituted by lower level sub-networks, each of them must produce the outcomes required by the functionality of the higher network, no matter what modes of operation they use to produce it. The way in which this hierarchical control is exerted is by means of information feedback control as we will see in §5.

During the constitution of an equivalence class, and in general when top-down causation through information control is at work, one can expect the selection of fewer different types of structural motifs belonging to an equivalence class of larger networks than one would normally expect given the mere combination of lower level elements (even smaller networks). This is due to the fact that the information of the network of higher complexity strongly affects the nature of information of the lower nodes, reducing the number of possibilities as the number of constraints within the larger system increases. The best analogy we have for RNA is that of a language. With RNA, it is possible to speak many 'words' that have the same 'meaning', as they belong to an equivalence class. When these words are linked to one another, they define a 'sentence' that conveys a greater meaning than its more basic components. Now the question might be whether it is the information of the words at the lower level which is affecting the information carried by the whole sentence, or the opposite. In fact, one can realize that what may be critical for the information carried by the whole sentence is the syntax that connects the words together, rather than the words themselves.

5. INFORMATION CONTROL

5.1. What is information control?

Since we are focusing on top-down causation by information control, we have now to see how information control should be characterized. We have already defined the two concepts that are basic for any information acquisition process, namely mutual information and information selection.

One of the biggest misunderstandings in information theory is to have taken [Shannon's \(1948\)](#) theory of communication (in the context of controlled transmission) as a general theory of information. In such a theory, centred on signal/noise discrimination, the message is already selected and well defined from the start, since from the start the selection operation among several alternative states (or bits) has already occurred (at the input or sender), and the problem here is only to faithfully transmit or further process, in the presence of disturbances, the sequence of bits that has been selected ([Auletta 2008a](#)). On the contrary, a true information theory (as was [Wiener's \(1948\)](#) original aim) starts with an input as a *source of variety* and has the selection only *at the end* of the information processing or exchanging. In other words, a message here is only the message selected by the *receiver*. As a matter of fact, *any* information reception will be subject to the original variety, in addition to the consequences of disturbance, dispersion or even of practical needs, and use of any of this information, at the most elementary level, already constitutes information selection. This is momentous for biological systems, since they are confronted with an environment that represents sources of uncertainty, and for this reason do not have control from the start of the string of bits that has been sent. Even inside a single cell we have such a problem, due to the modularization of the different subsystems. The control must here somehow be exerted while having only a limited pool of resources.

Let us state the problem in this way. According to the traditional information (communication) theory, the main problem is reliability, understood as the matching between input and output. However, in biological phenomena we are much more interested in situations in which the receiver does not have full control over the input and is therefore *forced to guess the nature of the input by taking the received partial information as a sign of it* (revealing its nature). This provides the more basic condition for equivalence classes, in that it is possible to take a partial input as a sign of many (possible) different entire input situations and to regard different inputs as equivalent under a certain point of view.

At any biological level, the receiver is, in general, flooded with incoming data and has to separate *background data* (important but constant) and *noise* (irrelevant data) from *relevant information* (data that are needed for some purpose). Therefore, *information control* consists of information selection—often involving a guess—from a certain point of view and this represents the goal of the system. For instance, a bacterium searching for an energy source may use a specific temperature gradient (the received information) as a

Figure 5. The basic feedback control scheme. (a) In living beings, feedback control underlies homeostatic processes (maintaining blood pressure, body temperature, etc.). The comparator determines the difference between the system state and the goal, and sends an error signal, activating the controller to correct the error. This is the way that goals (abstract variables) control dynamic causes, thus becoming causally effective. The goals in simple biological systems are genetically determined (unchanging through the organism's life and constant across a species). (b) The application of this framework to present day cellular systems. We show the biopolymers involved in each step.

sign (see §5.3) of this source. Obviously, many different temperature distributions (within a certain window) will fit and therefore allow it to reach a certain source, which is the goal. Moreover, any source that fits some general criterion will be good and therefore pertains to the equivalence class established by the goal of acquiring energy. We now need to state how goals and feedback control are linked. Information control via feedback is not the only way to have control via information but it plays a fundamental role in living systems: it is indeed involved, for example, in any homeostatic process that a living system must perform to survive.

5.2. Goals and feedback control

We can speak of top-down causation by information control when there are two elements that cannot be reduced to any low-level explanation: (i) the *formal structure* determining the feedback control loop (the formal cause, in our language) and (ii) the *goal*. These two elements represent the way the equivalence class is built in and controlled from above. To exercise information control, we need to join a system (the controller) with another system (the performer), see [figure 5](#). However, not all forms of connection, but only those satisfying the following requirements, will work:

- (i) the performer has to execute an operation in order to deploy the function needed by the controller, and
- (ii) the controller needs to be able to verify step by step whether the function is actually deployed to the required degree.

The requirement (i) is strictly linked to the fact that the controller has an inbuilt goal to reach; and the requirement (ii) is fulfilled by a feedback circuit (and an inbuilt comparator). It seems that mechanical devices, like a thermostat, are able to implement information control without any intervention of biological elements

and in purely mechanical terms. This is, however, an erroneous point of view, since such devices have been built by humans to act in a certain way. Therefore, the functional element (and the goal) is already inbuilt.

Feedback control (see again figure 5) is therefore the way information is causally effective in the biological world (Milsum 1966; Calow 1976; Ellis 2006a). Feedback structuring is necessary in making the information useful for attaining goals, otherwise the controller cannot recognize whether it has succeeded in obtaining a desired outcome. In this way, the controller is able to use a specific operation by a performing system as an element of a functional equivalence class.

At lower biological levels, the goal is somehow built into the controller and does not itself need to be 'chosen' (the situation is different when consciousness is involved). The goal state (relative to a certain specific parameter) is not a need *for the controller itself*, an entity which is a part, a *sub-system*, of the whole biological system. Moreover, that the goal state represents a need for the whole biological system does not mean this is an end *in itself*, since in general the exigency of keeping the system state coincident with the goal state depends on the latter being crucial for *other* operations or functions, in turn indispensable for the whole system. An example of information control is the expression or the repression of segments of DNA. Indeed, this control procedure, done on strings containing information and through activators and repressors that carry and instantiate instructions, is highly contextual and depends on the goals that the organism pursues in different time windows and states of its developmental or metabolic activity.

When an operation is actually performed, the fact that the outcome satisfies the goal (so that the comparator no longer sends error signals) can be seen as evidence for the operation pertaining to the equivalence class defined by the goal itself (this is the reason why functional equivalence class is not only an epistemic category). The goal can be satisfied by several slightly different outcomes, all of them falling into a well-defined range of tolerance. In order for an operation to be acknowledged as a member of a functional equivalence class, it is sufficient to individuate a specific feature that is tightly bound with the outcome achievement. Therefore, when we speak of equivalence classes at a biological level (functional equivalence classes), we enter the domain of signs (see §5.3). That the goal defines the equivalence class related to it is reminiscent of the concept of equifinality, a type of convergence, introduced by von Bertalanffy (1969) in the framework of general system theory.

5.3. Signs and equivalence classes

The connection between information control (and top-down causation) and equivalence classes is exactly represented by the goals and the related issue of signs.

In our usage, a *sign* is something that stands for something else in a certain context, or for a certain goal (Peirce 1931–1958, CP 2.228 and 1.540; Auletta 2007, 2008b), and this is what establishes equivalence classes in a context of information control. Signs convey

Figure 6. Information and signs. Input information, as a source of variety, starts an informational process that is concluded when information selection is accomplished. When this selected element is taken to say something about the input information, we have a semiotic relation and say that it is a sign of the input. This is evidently true when this element is tightly connected to that complex of initial information; something, however, can also be a sign of things that are not obvious consequences of the input information, for instance when certain items are taken to be a sign of the needed element (able to satisfy the goal).

information that has to be recognized and this occurs even in biological systems without consciousness (Hoffmeyer 1996; Deacon 1997).

As we have seen, in any information exchange we have selection *at the end*, not at the start. Signs, in the most elementary case, take this output selection as saying something *about the input*, so that the receiver starts a new information process aiming at the source (inverting somehow the ordinary flow of information from the source to the receiver, figure 6). There are many examples of the use of signs in biology. Among the most remarkable ones are the so-called *affordances* (Gibson 1966, 1979), when an animal takes a physical input (a smell) as the sign of something that is fundamental for survival (food).

With respect to equivalence classes, the outcome of the operation performed in a certain biological context may be considered as a sign of the function being deployed successfully (or not), and hence that the operation really is (or not) a member of the equivalence class. In other words, of all the input information entrained in the physical–chemical properties of the molecule or of the interaction under consideration, a *single* feature is selected and taken as the sign of the function required by the inbuilt goals of the whole system (as we have said, it is a part–whole relation). This sign may be the outcome of the operation, or any of its features reliably associated with the outcome. Let us call this sign a *mark* of the operation. Systems able to perform a certain operation for a certain goal are acknowledged through a specific mark and controlled in their mode of operation through such a mark. Thus the informational control of an operation (see §5.2) is a genuine semiotic process, since the control instance needs to catch and select specific information as a sign of the fact that things are going in the right or the wrong way.

5.4. Functional selection and top-down causation

In many cases, new information is acquired through functional selection—and therefore specialization—of determined operations at the lower level. Strong

Figure 7. Top-down causation. The relations between the elements involved in our theoretical scheme. The arrows show the conditions that contribute to the establishment of an operation.

evidence for this is given by mutation and heritage of epigenetic mechanisms, like the restoration of native folding of single-stranded DNA sequences through reverse mutations (Shepherd *et al.* 2006). When such a specialization occurs, it becomes truly useful for attaining a certain higher level goal, once it is linked to the previous set of controllable information of the biological system through a feedback loop, eventually establishing and maintaining a new functionality. Such a process is also known, at an evolutionary time scale, as *exaptation* and consists, in general, of the use of adaptive traits in a new context different from that from which they have originally arisen (Gould & Verba 1982). This situation also enables *biological systems to undergo adaptive selection at a functional level*.

An excellent example is the transition from an RNA-based genome to the DNA-based genome (Jaeger & Westhof 2001). RNA shows two fundamental chemical instabilities: (i) instability of the ribose–phosphate backbone that can hydrolyze inherently to the presence of a hydroxyl group at the level of the sugar moiety, and (ii) instability of cytosine that can hydrolyze into uracil. At this level, these instabilities prevent the genome from growing too much due to an increasing loss of information through these chemical instabilities and this represents a limitation of information coding. To circumvent this major problem, two new metabolic pathways have been developed through adaptive processes. One corresponds

to the formation of deoxynucleotides from nucleotides, removing the first instability. The other one is much more remarkable and consists of the transformation of uracils into thymines by addition of one carbon methyl to the base. The last change allows a mechanism to repair a damaged or mutated genome to emerge.

This key step in evolution provides a better vehicle of information coding (DNA) as well as allowing RNA to be more dedicated to operative tasks rather than to long-term storage functions within the cell. It also provides the cell with a better control over the expression of its various functional components (DNA, proteins and RNA) by allowing modulation and better timing control. In other words, this increase in information control corresponds to an increase of the modularity of the biological system by specialization of its functional tasks. Obviously, such a *specialization* may work only if it happens in a suitable network of information control.

5.5. Our whole conceptual framework

Top-down causation by information control occurs thanks to the connection between equivalence classes and information control. We have top-down causation by information control when, once an equivalence class has been established, the information selection defining the operations that are the elements of the class is

Figure 8. Operations and functional equivalence. (a–c) Different operations connected with differences in networks that are notwithstanding functionally equivalent (adapted with permission from Massire *et al.* 1997). a(i)–c(i), RNase P RNA of type A; a(ii)–c(ii), RNase P RNA of type B.

conserved, thanks to modularity, despite the variability of lower level variables. In this case, the feedback control circuits produce reliable responses to higher level information (Ellis 2006b, 2008), allowing equivalence classes of lower level operations that give the same higher level response for a certain goal. As equivalence classes are abstract configurations rather than physical states and are (through the lower level modality of operations) the causally effective higher level entities in this context, we have top-down causation that cannot be reduced to any specific bottom-up or same-level causation. The system as a whole is illustrated in figure 7.

6. BIOLOGICAL MODELLING

6.1. A model of an equivalence class

The concept of equivalence class is readily available at the level of stable RNA biopolymers. Sequence and structural analysis of natural RNA molecules such as large self-folding catalytic RNAs (group I and group II introns, RNase P RNA, ribosomal RNAs) reveal that their conserved structural catalytic cores are often stabilized by peripheral modules or tertiary motifs that, despite their different local structures, contribute in a similar fashion to the final core assembly and stabilization (Jaeger *et al.* 1994, 1996; Westhof *et al.* 1996; Massire *et al.* 1997, 1998).

As shown in figure 8a, the two-dimensional structural networks of two different molecules of

RNase P RNA have a conserved structural core able to perform the same catalytic function that is the maturation of tRNAs (see also figures 2 and 3). They, however, differ in their peripheral regions, which are involved in the proper intra-molecular recognition of their two constitutive domains. As seen in figure 8b,c, this intra-molecular recognition function that allows spatially forming a tertiary contact between the helical element P9 and the helical element P1 is operated by two kinds of tertiary structural motifs that have different sequences and three-dimensional structures. For this reason, the modality of operations is also different. However, these two specific operations (and the related different tertiary motifs) are perfect examples of items belonging to the same equivalence class (RNA–RNA recognition). These components have been shown to be swappable within the structural context of a group I ribozyme, demonstrating that they are functionally equivalent (Jaeger *et al.* 1994).

6.2. Experiments on top-down causation

Notwithstanding the interesting results reported in §6.1, we wish to find stronger evidence for top-down causation by showing the establishment of equivalence classes by systems displaying information control. In §2, we have reported about an important experiment showing that, notwithstanding differences at the physical-chemical level, enzymes pertaining to

Figure 9. Demonstration of functional equivalence within the metabolic network of a unicellular organism. (a) See text. (b) The same function can be carried out by molecules with different operations. Operations a and b, which are carried out by RNA alone, are different due to the way they fold. Operation c takes advantage of a protein to help the RNA fold. Operation d is mediated by protein alone at the assembly and catalytic levels.

different organisms can be interchanged without losing their functionality. In our opinion, that experiment is a beautiful instance of the fact that different lower level operations are *informationally controlled from above* and constitute or may constitute functional equivalence classes.

For this reason, the experiment by Wegscheid *et al.* can be framed within a general model for testing top-down causation. Here (see figure 9a), one takes a unicellular organism X and identifies an operation a (we recall that an operation is a pathway of physical–chemical interactions) deploying a certain function in X (e.g. the maturation of tRNA). Then, one takes another organism Y deploying an alternative operation b and substitutes it for a in X. Then, one considers whether operation b is able to deploy the same function in the new context. If operation b is actually ‘accepted’ by organism X, then the role of functional equivalence class for that organism is proved to the extent to which the specific interaction pathway is disregarded and only the function is taken into account.

This means that the equivalence class in this case is established inside a *single* organism, suggesting that this happens through an information control feedback loop. Notice that transindividual equivalence classes (as in the experiments reported in §6.1) could not be considered as evidence for top-down causation but rather as a way the scientist classifies certain physical and chemical interactions. In fact, in this case what happens within two (or more) different organisms could be explained in purely mechanical and bottom-up terms in each organism separately. When, on the contrary, a single organism X is able to deploy a certain function either by performing operation a or b, this can be no longer understood in a purely mechanistic way, and the goal of the operation (as well as the piloting control instance) gets primary importance.

However, within the context of the Wegscheid experiment, it is important to stress that while type A and type B RNase P have significantly different

architectures, they are still evolutionarily related as their secondary and tertiary structures share a consensual catalytic site for binding and cleaving tRNA molecules. The two architectures both result from the evolutionary divergence of a common RNase P RNA ancestor. This is clearly established by extensive sequence comparative analysis and crystallographic analysis of the structure of the two classes of RNase P RNA (e.g. Altman 2007; Kirsebom 2007). Type A and type B RNase P enzymes have evolved through time to retain in their common structural core all the specific structural determinants for performing the same catalytic reaction of the maturation of pre-tRNAs. Therefore, a reductionist explanation such as the one proposed by Wegscheid *et al.* for their experiment cannot be completely ruled out.

The basic scheme of the type of experiment proposed is quite common in genetics, when a gene coding for a certain function is substituted by another similar one presenting some genetic variations.

Therefore, what we really want to perform is an experiment of substitution of molecules of the same function but with different modes of structural and functional operations, meaning different ways to perform the catalysis and recognize the substrate. This would show that two different evolutionarily unrelated molecular systems that share the same function by convergent evolution can substitute for one another. In fact, in living cells, there are numerous examples of such types of systems (e.g. DNA nucleases, aminoacyl-tRNA synthetases, self-splicing introns, self-cleaving ribozymes, etc.).

Therefore, a further refinement of the experiment (figure 9b) described above could be to eventually substitute a key cellular operation performed by a protein (like a nuclease activity) by one performed by an RNA. The functional RNA does not have to be of natural origin as it can eventually result from *in vitro* experiments (e.g. Jaeger 1997; Chworos *et al.* 2004; Chworos & Jaeger 2007). In this case, while the function could be the same (this is our prediction),

any level of organization below that one would be different, allowing demonstration of top-down causation by information control without ambiguity. Here, we assume that the organism level is higher than that of operations, and therefore anything that could be swapped, maintaining the same cell activity and its outcomes, is seen as an effect of top-down causation.

6.3. A research programme

In this paper we propose a new research programme, through which we try to experimentally establish top-down causation. It is obviously a very difficult issue and, probably, many experimental steps as well as new types of experiments will be necessary. However, we also believe that, given the conceptual framework developed here, it is a possible enterprise.

It is always difficult to positively prove a result, so one may wonder whether it is possible to positively prove top-down causation. In complex cases, one could always object that there will be at some time a purely molecular (and bottom-up) explanation of the eventual findings that we are aiming at. However, showing that equivalence classes are constituted at the most basic biomolecular level represents a strong counterexample to all bottom-up and same-level explanations, strongly suggesting top-down causation as a fruitful way to understand these results. It will also help to explain issues regarding emergence and related evolutionary aspects which constitute the major part of our future programme. It should be also clear that a possible failure of our research does not necessarily imply a direct confutation of top-down causation. Other ways to search for it, both at the microbiological level and at higher levels of complexity, are likely to be found.

We wish to express our warmest gratitude to Ivan Colagé and Paolo D'Ambrosio, two doctoral students of the Pontifical Gregorian University, who actively participated in the discussion at the workshop held in Rome on 24–28 September, which made this paper possible. We thank Horst Klump, of the University of Cape Town, for the example of native folding of DNA sequences, and Bill Stoeger, of Specola Vaticana, for helpful comments. The paper has been significantly improved through constructive comments of our referees. The work was funded in part through a grant from the Center for Theology and the Natural Sciences, Berkeley, California (www.ctns.org).

REFERENCES

- Altman, S. 2007 A view of RNase P. *Mol. Biosyst.* **3**, 604–607. (doi:10.1039/b707850c)
- Auletta, G. 2005 Quantum information as a general paradigm. *Found. Phys.* **35**, 787–815. (doi:10.1007/s10701-005-4565-6)
- Auletta, G. 2006 The problem of information. In *Proc. I Workshop on the Relationships Between Science and Philosophy* (ed. G. Auletta), pp. 109–127. Vatican City, Italy: Libreria Editrice Vaticana.
- Auletta, G. 2007 Information, semiotics, and symbolic systems. *Semiotica* **166**, 358–376.
- Auletta, G. 2008a Biological systems integrating information and entropic fluxes. In *The controversial relationships between science and philosophy: a critical assessment*, pp. 13–22. Vatican City, Italy: Libreria Editrice Vaticana.
- Auletta, G. 2008b Autarchy and openness in living systems. In *The controversial relationships between science and philosophy: a critical assessment*, pp. 23–38. Vatican City, Italy: Libreria Editrice Vaticana.
- Auletta, G. 2008c How many causes there are? *XXI Secolo* **7**.
- Barabasi, A.-L. & Oltvai, Z. N. 2004 Network biology: understanding the cell's functional organisation. *Nat. Rev. Genet.* **5**, 101–114. (doi:10.1038/nrg1272)
- Bishop, R. C. & Atmanspacher, H. 2006 Contextual emergence in the description of properties. *Found. Phys.* **36**, 1753–1777. (doi:10.1007/s10701-006-9082-8)
- Booch, G. 1994 *Object oriented analysis and design with applications*. Reading, MA: Addison Wesley.
- Calow, P. 1976 *Biological machines: a cybernetic approach to life*. London, UK: Edward Arnold.
- Campbell, D. T. 1974 Downward causation. In *Studies in the philosophy of biology: reduction and related problems* (eds F. J. Ayala & T. Dobzhansky), pp. 179–186. Berkeley, CA: University of California Press.
- Chworos, A. & Jaeger, L. 2007 Nucleic acid foldamers: design, engineering and selection of programmable bio-materials with recognition, catalytic and self-assembly properties. In *Foldamers: structure, properties, and applications* (eds S. H. Hecht & I. Huc), pp. 291–330. Weinheim, Germany: Wiley-VCH.
- Chworos, A., Severcan, I., Koyfman, A. Y., Wienkam, P., Oroudjev, E., Hansma, H. G. & Jaeger, L. 2004 Building programmable jigsaw puzzles with RNA. *Science* **306**, 2068–2072. (doi:10.1126/science.1104686)
- Clayton, P. & Davies, P. C. W. (eds) 2006 *The re-emergence of emergence*. Oxford, UK: Oxford University Press.
- Collins, S. R. et al. 2007 Functional dissection of protein complexes involved in yeast chromosome biology using a genetic interaction map. *Nature* **446**, 806–810. (doi:10.1038/nature05649)
- Crick, F. 1994 *The astonishing hypothesis. The scientific search for the soul*. New York, NY: Scribner. (Reprinted by Touchstone 1995.)
- Deacon, T. W. 1997 *The symbolic species: the co-evolution of language and the brain*. New York, NY: W. W. Norton.
- Deacon T.W. 2006 Emergence: the hole at the wheel's hub. In *The re-emergence of emergence* (eds P. Clayton & P. Davies), pp. 111–150. Oxford UK: Oxford University Press.
- Dowe, P. 1992 Wesley Salmon process theory of causality and the conserved quantity theory. *Philos. Sci.* **59**, 195–216. (doi:10.1086/289662)
- Dowe, P. 1995 Causality and conserved quantity theory: a reply to Salmon. *Philos. Sci.* **62**, 321–333. (doi:10.1086/289859)
- Ellis, G. F. R. 2006a Physics and the real world. *Found. Phys.* **36**, 227–262. (<http://www.mth.uct.ac.za/~ellis/realworld.pdf>)
- Ellis, G. F. R. 2006b On the nature of emergent reality. In *The re-emergence of emergence* (eds P. Clayton & P. C. W. Davies), pp. 79–107. Oxford, UK: Oxford University Press.
- Ellis, G. F. R. 2008 On the nature of causation in complex systems. *Trans. R. Soc. S. Afr.* **63**, 1–16.
- Fair, D. 1979 Causation and the flow of energy. *Erkenntnis* **14**, 219–250. (doi:10.1007/BF00174894)
- Forster, A. C. & Church, G. M. 2006 Towards synthesis of a minimal cell. *Mol. Syst. Biol.* **2**, 45. (doi:10.1038/msb4100090)
- Gibson, J. J. 1966 *The senses considered as perceptual systems*. Boston, MA: Houghton Mifflin. (Reprinted by Greenwood Press 1983.)
- Gibson, J. J. 1979 *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin. (Reprinted by L. Erlbaum Associates 1986.)

- Gould, S. J. & Verba, E. 1982 Exaptation: a missing term in the science of form. *Paleobiology* **8**, 4–15.
- Hebb, D. O. 1949 *The organization of behaviour: a neuro-psychological theory*. New York, NY: Wiley. (Reprinted by Erlbaum 2002.)
- Hoffmeyer, J. 1996 *Signs of meaning in the Universe*. Bloomington, IN: Indiana University Press.
- Holland, J. 1997 *Emergence: from chaos to order*. Reading, MA: Addison-Wesley.
- Jaeger, L. 1997 The new world of ribozymes. *Curr. Opin. Struct. Biol.* **7**, 324–335. (doi:10.1016/S0959-440X(97)80047-4)
- Jaeger, L. & Westhof, E. 2001 Le monde de l'ARN. *L'Astronomie* **115**, 284–288.
- Jaeger, L., Michel, F. & Westhof, E. 1994 Involvement of a GNRA tetraloop in long-range RNA tertiary interactions. *J. Mol. Biol.* **236**, 1271–1276. (doi:10.1016/0022-2836(94)90055-8)
- Jaeger, L., Westhof, E. & Michel, F. 1996 Function of a pseudoknot in the suppression of an alternative splicing event in a group I intron. *Biochimie* **78**, 466–473. (doi:10.1016/0300-9084(96)84753-5)
- Kirsebom, L. A. 2007 RNase P RNA mediated cleavage: substrate recognition and catalysis. *Biochimie* **89**, 1183–1194. (doi:10.1016/j.biochi.2007.05.009)
- Küppers, B.-O. 1990 *Information and the origin of life*. Cambridge, MA: MIT Press.
- Lashley, K. S. 1942 The problem of cerebral organization in vision. *Biol. Symp.* **7**, 301–322.
- Lehn, J.-M. 2004 Supramolecular chemistry: from molecular information towards self-organization and complex matter. *Rep. Prog. Phys.* **67**, 249–265. (doi:10.1088/0034-4885/67/3/R02)
- Lescoute, A. & Westhof, E. 2006 Topology of three-way junctions in folded RNAs. *RNA* **12**, 83–93. (doi:10.1261/rna.2208106)
- Massire, C., Jaeger, L. & Westhof, E. 1997 Phylogenetic evidence for a new tertiary interaction in bacterial RNase P RNAs. *RNA* **3**, 553–556.
- Massire, C., Jaeger, L. & Westhof, E. 1998 Derivation of the three-dimensional architecture of bacterial ribonuclease P RNAs from comparative sequence analysis. *J. Mol. Biol.* **279**, 773–793. (doi:10.1006/jmbi.1998.1797)
- Milsum, J. H. 1966 *Biological control system analysis*. New York, NY: McGraw Hill.
- Murphy, N. & Stoeger, W. R. 2007 *Evolution and emergence: systems, organisms, persons*. Oxford, UK: Oxford University Press.
- Pearl, J. 1998 Graphs, causality, and structural equation models. *Sociol. Methods Res.* **27**, 226–284. (ftp.cs.ucla.edu/pub/stat_ser/R253.ps)
- Pearl, J. 2000 *Causality: models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.
- Peirce, C. S. 1931–1958 [CP] *The collected papers of Charles S. Peirce*, vols. I–VI (eds C. Hartshorne & P. Weiss). Cambridge, MA: Harvard University Press, 1931–1935, vols. VII and VIII (ed. A. W. Burks). Cambridge, MA: Harvard University Press, 1958.
- Rasmussen, S., Chen, L., Deamer, D., Krakauer, D. C., Packard, N. H., Stadler, P. F. & Bedau, M. A. 2004 Transitions from nonliving to living matter. *Science* **303**, 963. (doi:10.1126/science.1093669)
- Roederer, J. G. 2005 *Information and its role in nature*. Berlin, Germany: Springer.
- Salmon, W. 1984 *Scientific explanation and the causal structure of the world*. Princeton, NJ: Princeton University Press.
- Salmon, W. 1994 Causality without counterfactuals. *Philos. Sci.* **61**, 297–312. (doi:10.1086/289801)
- Schmid, E. M. & McMahon, H. T. 2007 Integrating molecular and network biology to decode endocytosis. *Nature* **448**, 883–888. (doi:10.1038/nature06031)
- Shannon, C. E. 1948 A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423. See also pp. 623–656.
- Shepherd, D. N., Martin, D. P., Varsani, A., Thomson, J. A., Rybicki, E. P. & Klump, H. H. 2006 Restoration of native folding of single stranded DNA sequences through reverse mutations. An indication of a new epigenetic mechanism. *Arch. Biochem. Biophys.* **453**, 108–122. (doi:10.1016/j.abb.2005.12.009)
- Simon, H. A. 1962 The architecture of complexity. *Proc. Am. Philos. Soc.* **106**, 467–482.
- von Bertalanffy, L. 1969 *General system theory: foundations, development, applications*. New York, NY: G. Braziller.
- Wegscheid, B., Condon, C. & Hartmann, R. K. 2006 Type A and B RNase P RNAs are interchangeable *in vivo* despite substantial biophysical differences. *EMBO Rep.* **7**, 411–417.
- Westhof, E., Masquida, B. & Jaeger, L. 1996 RNA tectonics: towards RNA design. *Fold. Des.* **1**, 78–88. (doi:10.1016/S1359-0278(96)00037-5)
- Wiener, N. 1948 *Cybernetics, or control and communication in the animal and in the machine*, 2nd edn. Cambridge, MA: MIT Press. (Reprinted 1961, 1965.)