

W.D. Christensen and C.A. Hooker
Philosophy, University of Newcastl

Organisation, Evolution and Cognition: Beyond Campbell's Evolutionary Epistemology

Abstract

Donald Campbell has long advocated a naturalist epistemology based on a general selection theory, with the scope of knowledge restricted to vicarious adaptive processes. But being a vicariant is problematic because it involves an unexplained epistemic relation. We argue that this relation is to be explicated organisationally in terms of the regulation of behaviour and internal state by the vicariant, but that Campbell's selectionist account can give no satisfactory account of it because it is opaque to organisation. We show how organisational constraints and capacities are crucial to understanding both evolution and cognition and conclude with a proposal for an enriched, generalised model of evolutionary epistemology that places high-order regulatory organisation at the centre.

I. Campbell's theory of knowledge.

Donald Campbell develops a naturalist theory of epistemology based on a general selection theory which posits three major components to selectionist processes: the production of variants (V), selection across these variants (S), and the retention of those variants which survive the selection process (R). The major intuition underlying Campbell's theory is that knowledge emerges through, and only through, VSR adaptation to the environment (e.g. Campbell 1974). In this respect the epistemic sophistication of a scientist compared with his virus-type ancestor, or of science compared with their virus-type ancestral community, simply reflects the cumulative inductive achievements of millennia of Darwinian evolution (blind-VSR). In this respect Campbell proposes that there occur intrasystemic processes which function to shortcut evolution by vicariously anticipating characteristic environmental conditions (so either avoiding or encouraging them), and the knowledge of these processes is derived from the VSR processes which produced them, and these vicarious processes themselves operate according to VSR principles. More recently Campbell importantly qualified this picture by restricting the scope of knowledge to just these vicarious processes (Campbell 1997). This leaves us with the following account of knowledge. State I of system S is knowledge of environmental condition C iff: (a) I is correlated with C because of a selective history, and (b) S treats I as a vicariant for C.

The qualification restricting knowledge to vicariant processes is important because it rules out many problematic cases and potential counterexamples. For example, contemporary viruses have apparently been no less structured by VSR processes than has the scientist, but they are no more epistemic than their ancient ancestors. This is in sharp contrast with the lineage of the scientist, which is marked by a general tendency towards increasing epistemic sophistication realised through complexly organised neural structure. Since a VSR process is common to both cases, VSR structuring in itself is at best a necessary feature of epistemic systems, and the virus case makes it look a relatively uninformative requirement.

This counterexample is ruled out by Campbell's restriction-to-vicariance condition because only in the case of the scientist's lineage has there appeared vicarious processes, in the form of a complex,

neurally supported functional organisation. This organisation forms complexly nested conditional response structures and we suppose with Campbell that at least the genetically inheritable component of these arose through VSR processes. Thus we have VSR processes for the refinement and conditionalising of the products of earlier VSR processes, i.e. second and higher order VSR processes. So both the contemporary virus and the scientist have experienced first-order VSR structuring, but only the scientist has experienced second- and higher-ordered VSR structuring, and it is this high-order VSR structuring which is constitutive of the scientist's massively enhanced epistemic capacity. Roughly, VSR epistemology involves interaction with the environment which selects from amongst competing beliefs, where beliefs are essentially states which are vicariants for environmental conditions. Campbell's extension of VSR epistemology to science per se involves showing that science as a cultural entity creates and sustains vicarious VSR processes which allow the referents of scientific beliefs to play a role in selecting the belief structures of scientists.

The distinction between a system feature which is merely adaptive and a feature which is a vicariant plays a critical role in Campbell's theory, but it is itself somewhat problematic inasmuch as being a vicariant involves an unexplained epistemic relation. That S uses I as a vicariant for C means roughly that S treats I as an indicator of C, however what the indication relation comes to is unclear. It cannot simply mean that I is correlated with C because any system feature which is an adaptation carries some degree of mutual information with the environment (though see Christensen/Collier/Hooker 1997 section V.2), and consequently would meet the criterion irrespective of its epistemic significance. The least question begging, and from our perspective most interesting, interpretation of the indication relation we can make is that I indicates C for S if I plays a role in regulating S's processes so that they are appropriate to C. This means not only that I has been selected for correlation with C, but that I is embedded within a more general regulatory context in which S has a control relationship with C.¹

However, although a regulatory interpretation of vicariance saves Campbell's VSR epistemology from being question begging, his relative neglect of it raises several further problems for him. (1) Campbell repeatedly emphasises that 'competence of reference' is the hallmark of a vicariant state qua knowledge, where he seems to understand this condition in a purely correspondence manner. However competent process regulation involves much more than competence of reference in this sense, which is useless in itself unless the system can make appropriate use of the vicariant. The extent to which I is a vicariant of C for S depends as much on the systemic context of I - in particular on the way C interacts with S and S's goals - as it does on I's referential competence. Indeed, treating competence of reference as the primary epistemic relation is misleading: to function as a regulatory vicariant for C, I must modulate S's processes in ways that are appropriate to S's interaction with C, but I needs only limited competence of reference to do this, it simply needs to control certain aspects of S's interaction with C. The mosquito flying up the CO₂ gradient does not need reference to it even as an extended stream, it only needs to control its flight direction by its local gradient. Campbell's theory is somewhat vague concerning the epistemic relations involved in vicariance, because it almost entirely neglects the contextual regulatory aspects which must be understood to properly account for vicariance.² (2) Selection theory is fundamentally unable to address regulatory issues. Despite widespread attempts to the contrary, it is simply the wrong kind of theory. Consequently, whatever the role of VSR processes may be in knowledge development, contra Campbell's central thesis they do not represent a sufficient condition for its occurrence (cf. also Hooker 1995). We now turn to

elaborating this point.

II. Why selection theory is unsuited as a foundation for natural epistemology.

The basic problem with selection theory as a theory of natural epistemology is that it is radically impoverished as an explanatory model of organised systems, but understanding organisation is central to epistemology. Selection theory takes a narrow slice through a very complex phenomenon, and one partially inappropriate for epistemic purposes. In particular, it washes out or glosses over almost all of the actual dynamics of the system being modeled, being focused just on frequencies of outcomes. This is not a problem if the theory nevertheless picks out the right properties for characterising the phenomenally/dynamically significant dimensionality of the system - as it does if the primary objective is just to understand the dynamics of population statistics. However we contend that in the case of selectionist models of organised systems many of the dynamical features which selection theory washes out play important roles in understanding the epistemically significant features of complex organised systems.

Standard evolution theory takes as its primary unit the population, and measures changes in gene frequency within the population over time. Much of the usefulness of selectionist models derives from being able to choose conditions such that the adaptiveness, or mutual information, of a particular gene is determined by its relative change in frequency as compared with other genes within the population. This occurs when the population is treated as a decomposable, near-to-equilibrium system. These linearising conditions are often only implicit in the models, but only under such conditions can outcome distributions be treated as independent of dynamical path from their starting values (e.g. as in the Hardy-Weinberg law). Moreover, these assumptions are central to the explanatory capacity of the theory: If the population is not decomposable into quasi-independent systems (organisms, genes), because it displays holistic dynamical constraints, then relative frequency measures may not be well defined. There are many systems of this kind, and plausibly organisms and social systems - the epistemic systems in question - are amongst them (see below). Further, even if the system is decomposable or nearly so but the population is not near-to-equilibrium with its environment, then relative change in gene frequency cannot be assumed to be a good measure of adaptiveness because, beside correlation with the environment, path-dependency effects (e.g. founder effects in migration) will also play a role in determining frequency distribution. Note that we are not claiming that there cannot be nonlinear population-genetic models in these latter cases, just that in these cases relative frequency distribution will not directly measure adaptiveness.

However, just noting the limitations in the applicability of selectionist models is only part of showing that they represent an unsuitable foundation for natural epistemology - more important are the reasons for these limitations. The fundamental problem with selection theory derives from its opacity to dynamical and organisational factors. Although a selection model compares the relative frequency changes of stable traits over time, such a model provides no insight into the dynamical processes by which the frequencies evolve. Moreover, by only comparing intra-population differences, the contribution to adaptiveness of the organisation of the members of the population (e.g. individual organisms or social organisations) drops out of the picture - in effect, it is suppressed by the high level abstracted structure of selectionist models. However it is precisely such dynamical and organisational information which we need if we are to be able to make distinctions between various sorts of adaptive processes, e.g. between gene-based change of fixed first-order traits to improve

environmental fit and high-order modification of adaptable traits in organism-learned fit.

The problem cannot be rectified simply by extending selection theory to intra-organismic or social processes. Although any system with stable traits can have a black-box frequency counting model imposed on it (and some aspects of both organism and social structure may be susceptible to this, though this not the case generally) the basic explanatory gap remains - there is no account of process organisation. Organisms and social systems are marked by hierarchical, holistic organisation where such modularity as occurs (cells, organs, institutions) is heavily constrained by global functional organisation. These kinds of systems are not dynamically decomposable because: (a) their components are highly interdependent, and (b) their dynamics tends to amplify this holistic organisation. Consequently they can only be adequately modelled by a theory which explicitly treats dynamical and organisational factors.

If we are to successfully model knowledge development as an adaptive process - the root intuition behind Campbell's evolutionary epistemology - we must distinguish amongst different kinds of adaptive processes. As Campbell is all too aware, features of social systems (such as religious ideas) may be adaptive without being epistemically reliable. On the other hand epistemic reliability does seem to be a strong feature (if not the only one) of the adaptiveness of scientific ideas. Clearly there must be some difference between the adaptive processes of the respective systems, but as we have seen, although selectionist models may in certain circumstances be sensitive to differences in adaptiveness, they are not sensitive to the underlying reasons for these differences. Consequently selection theory cannot distinguish between adaptive processes which are distinctively epistemic and processes which are adaptive for other reasons. A theory of vicariant processes may well be an important step towards understanding distinctively epistemic adaptive processes (at least in the control form we have given to indication here), but it cannot be developed from within selection theory. We now provide a richer framework within which to develop these ideas.

III. Organisation, endogenous control and the generation of knowledge.

The problems with Campbellian vicariance are inherent in his attempt to generalise selection theory to organised systems, where the principled inadequacy to organisation of correlation (the only outcome selection can distinguish) defeats epistemic insight. The problem is manifest at a number of points in his work, resulting, e.g., in a chronic inability to discriminate in a principled way between science and other cultural systems, and a misguided account of the constraints associated with embodiedness: so-called 'vehicular' and 'co-selection' constraints on competence of reference. To deal with the latter issue first, organised systems characteristically face global constraints, and these constraints may function both in an inhibitory fashion, by ruling out certain otherwise available complexions of the system, and in a functionally amplifying fashion by making available certain capacities which the system would be otherwise unable to achieve. We refer to the latter as enabling constraints. For example, the cell membrane is a global enabling constraint with respect to intra-cellular biochemical organisation because it provides a quasi-isolated environment within which critical parameters such as ionic concentrations, enzyme activity etc., may be maintained and organised so as to perform useful work and without which they would collapse. Campbell, however, models all such constraints as limitations on epistemic capacity, e.g. characterising 'vehicular' constraints in terms of resolution and colouration limitations on reference capacity (1997, section 4.2.2); there is no hint that some constraints may play an enabling function. Similarly, Campbell

struggles to distinguish science from other types of cultural systems, resulting in the recognition that science like all cultural systems faces 'tribal' constraints. The problem derives from the fact that VSR structuring and embodiment constraints are common to both scientific and non-scientific cultural systems - the differences will lie in organisational features of science which play an enabling role in knowledge development.

For example, both the immune system (IS) and central nervous system (CNS) have VSR functional characteristics (as well as origins), but they have very different organisational - and epistemic - characteristics. If the IS is properly characterised as epistemic (and we think that in some respects it is), it is a very low-order epistemic system whose 'object reification' capacity is virtually nonexistent. The CNS, on the other hand, is an object reifier par excellence (though it is also much more than this), but not because it has 'more VSR' than the IS. Fundamentally, the CNS is capable of object reification because it has far greater control depth than the IS. If we consider the organisation of a primate retina, to take just one instance, we may observe that there are approximately 100 million photoreceptors which detect light signals which synapse onto a mere 1 million axons in the optic nerve (Churchland and Sejnowski 1992, p.148). Connections amongst receptors are highly organised, showing both mutual inhibition and activation. In other words, the behaviour of individual retinal neurons is heavily constrained by the local and global organisation of the retina, and it is precisely these constraints (together with those of the visual cortex) which enable the general pattern recognition capacity of the primate visual system.³ The CNS also has more VSR functional structuring than the IS, but this is parasitic upon its greater organisational depth, and it is the increased control depth in particular which is the critical enabling constraint which permits object reification. This is why at the outset we recast Campbell's notion of vicariance (as indication) in terms of organised control depth. We now turn to providing a wider framework within which to discuss the role of organised control depth in epistemic systems.

We begin by considering the structure of VSR processes themselves within organised systems. Campbell's account tells us nothing significant about this; as we have seen, it goes no further than to say that the system (including its organisation) emerges from sequences of VSR processes working on preceding VSR processes. That this is clearly inadequate to understanding is shown by the manner in which the V, S and R processes themselves are modified as organised control depth increases. The most general feature of this is increase in endogenous control of V, S and R.

In the standard neo-Darwinian conception, the process of Darwinian evolution is representable by a VSR process in which all three components are uncontrollable, they are not regulated by any controlled organisation. Variation, pure genetic mutation, is not controllable because it is random. Selection, here either death of, or reproduction of, the phenotype, is uncontrollable because it is a deterministic, exogenous outcome of environmental interaction. Retention appears as altered gene frequencies and is uncontrollable because it is a deterministic consequence of selection outcome, and in that sense continues to be exogenous to both phenotype and population. Call this process Simple Darwinian Evolution (SDE).

By contrast, when we consider the situation with the organisms of substantial phenotypic capacity because of their complex and deep endogenous organisation, we pass to a situation where each of V, S and R are under significant endogenous regulatory control. These organisms significantly

regulate (i) their own development, (in varying degrees) both their physiological and cognitive/behavioural development, (ii) the structure of their environment, both natural and social, and, through both, (iii) their behaviour, including their reproduction. This endogenously controlled shaping of endogenous capacity and exogenous social and natural environment is a capacity to (partially) endogenously control the VSR process. V: Endogenous control of effective variation per se is ubiquitous since individuals are not effectively presented for environmental selection unless they have had viable embryogeneses and accompanying social development trajectories, both of which are (partially) endogenously controlled. See also R below. S: Creatures actively modify their own environment, typically in proportion to their (socially amplified) phenotypic capacities, thereby modifying the selective force which they experience. Again, creatures that actively modify themselves and/or their community structures so as to increase survival skills effectively thereby modify the selection pressure that they experience, and this modification is under endogenous (and social) control. R: Controlled modification of inheritance, i.e. of the species memory of developmental factors, occurs in biased mutation directions and/or rates of, and in multiple reproduction of, fertilised ova, whether naturally or (now more commonly) through many socially regulated technologies. These processes, which also alter effective variation, are either directly or socially (hence indirectly) endogenously controlled. Call this process Complex Organised Darwinian Evolution (CODE).

From this perspective SDE is a degenerately simple form of CODE, omitting all the dynamically crucial organisational constraints. To understand evolutionary dynamics it is necessary to recognise the complex ways in which effective V, S and R come under increasing endogenous control as organised phenotypic and/or social capacities increase.

To best understand high depth control we need to remember that the basic function of VSR processes is to tune the processes of systems so as to better maintain their autonomy. The most basic system processes are anticipative feedforward processes by which the system anticipatively responds to conditions (in itself or the environment) which it detects. There can be higher order processes which modify those underlying them and which are either directly anticipative feedforward or indirectly anticipative error-correcting feedback. High depth control involves explicitly regulating the structure of underlying processes rather than merely randomly reshaping them when they fail. The former is a much more powerful adaptive capacity than is the latter, which is merely running an uncontrolled VSR process, because it can utilise higher order (more subtly correlated) properties of its environment related to its own organisational needs (themselves now extracted as increasingly high order features). The key to this kind of organisation is the ability to extract a control parameter from endogenous interaction outcomes (e.g. from environmental signals received) and feed it forward to re-organise the existing lower order feedforward processes so as to better achieve system autonomy. The process of parameter extraction and use, when directed to improvement of autonomy through error-corrective feedback, is self-directed anticipativeness and elsewhere we have characterised such processes at length (Christensen/Hooker 1997).

As control depth increases the processes become more recognisably cognitive because the addition of autonomy-supporting modifier processes is an increase in anticipative depth (direct and/or indirect) and the essence of cognitiveness is its anticipative character. When this is achieved by the construction of high order regulators we have a paradigm cognitive learning process, self-directed learning (SDL) and elsewhere (Christensen/Hooker 1997) we have shown how characteristic aspects

of these regulators take the form of predication, representation, emulation and conceptualisation.

Campbell himself saw that anticipative feedforward combined with error-corrective feedback supports powerful learning and response capabilities but missed the specific organisational form which connects these capacities to cognitive capacities proper. What then of the role of VSR processes in cognitive systems? Adaptive processes increase fixed or stored system mutual information with the environment and it can be said immediately that the most general form of all adaptive processes is VSR; all other processes occur at constant (or reduced) system information, showing internal variegation only. For variegation to be variation for a system it must be uncontrolled by, and promoted across, at least one cohesive system level; for this to increase system mutual information with the environment the promotion must be suitably selected; for the process to be operationally effective the information must be relevantly retained. However, these VSR processes may range from SDE and non-cognitive to high order cognitive CODE depending on the system resources available for adaptiveness in that respect. In addition, the execution of each self-directing anticipative process instantiates a controlled VSR process. Each such process initially presents a delimited range of possible parameter values for the extracted control parameter and its modification consequences (V control), specifies a range of signal kinds and values as those which can initiate the anticipative process (S control) and, following execution of modification, specifies any subsequent alteration of the process itself (R control). However, these VSR aspects appear only degenerately in adaptive processes below the organisation of self-directing processes.

High order controlled VSR processes are important to complex organised systems because: High order V control is a matter of focussing system resources to amplify variation within those bands of system variability which prove most productive in generating useful organisation. High order S control is a matter of using tuned filters sensitive to high order features as selective barriers. to allow the system much greater specificity in the kinds of organisational features to which it is sensitive. High order R control involves both increasing the organisation the system can retain and also controlling the way that organisation is then fed back into the adaptive processes of the system.

With this organisational framework in place we are ready to address Campbell's problem of satisfactorily characterising evolutionary epistemology, which we take to be a doctrine asserting some interesting relationship between the biological process of evolution and the development of science (and perhaps of cognition more generally). Briefly, according to the foregoing analysis, what makes science cognitive is its anticipative depth and SDL capacity. As we have now seen, this involves possessing high order controlled VSR processes. But we have also seen that the dynamics of the evolutionary process hinges on the roles of high order controlled VSR processes in like manner (crudely in proportion to the organisational depth of the phenotypes involved). This dynamical structure is what they share in common. (And, suitably modified, also what science shares in common with the internal anticipative processes of the individuals, the phenotypes, involved.) What then of the relationship of the two processes? In these circumstances we opt for the naturalist position that cognitive development generally is embedded within the evolutionary process, with science its current cutting edge (Hooker 1995). Now we can re-define evolutionary epistemology as the theory of the cognitive component of the CODE process. This is in contrast to all those 'standard' evolutionary epistemologies that take SDE as the basic process form, whereas it in fact provides only the degenerate 'outer form' of the learning process. The complex form of CODE, we believe, will

provide a rich framework for understanding the internal processes of science. In this paper however we confine ourselves to illustrating this with brief comments on controlled VSR processes in science.⁴

Endogenous V control: Theoretical proposals are not effectively presented for testing unless (i) they have already passed a stringent series of filters applied during their development, e.g. are clearly expressible within the currently available mathematical modelling and require no violation of basic principles (such as energy or PCT conservation), and (ii) unless they interrelate with the viable theories around them in an appropriately fruitful manner (e.g. yield workable models of experimental conditions) and do not contradict surrounding theories or make approximating ‘cuts’ through reality in very different ways from the current approximate models available. All this represents control of those theoretical variants which are presented for selection, control which is endogenous either to individuals or scientific groups.

Endogenous S control: A theory is tested in a situation where the key parameters are controllable, the features of epistemic interest are accessible to measurement and measured variables are sensitively monitored with specially designed instruments. This shaped, pre-selected environment profoundly shapes in turn the categories and ranges of signals which will be accepted as feedback from reality in testing theories. Finally, the raw experimental information, = instrumental behaviour, which results is subject to fastidious mathematical re-shaping before any scientific ‘data’ emerges. In this way, and in the ways scientists actively modify themselves through training, science exhibits a deep endogenous control of its own perceptual structures and hence of the selective force theories experience.

Endogenous R control: Scientists have elaborate, highly discriminatory institutionalised pathways that profoundly effect the way theories which pass tests are remembered or retained, i.e. are reproduced elsewhere throughout science. Successful theoretical proposals do not propagate universally but into highly selective contexts where their approximations and their particular control parameters and mathematical techniques can be fruitfully combined with others that exist at that active working site. Similarly, the development of a theory may be profoundly altered by the introduction of mathematical or experimental techniques from fields other than that in which it was originally proposed; the development of statistical mechanics, e.g., was profoundly altered by the development of mathematical statistics, especially the study of fluctuations, as was biology by the introduction of each new kind of microscope (optical, electron, tunnelling ...). Finally, some proposals, or even components of them, are subject to development of many alternative variants and their consequences carefully explored, while others are left in their original, unelaborated form. And of course scientists continually train their successor generations and re-train themselves in new techniques and in these ways also continually alter the composition of the inherited collection of skills and capacities and institutional designs which make up the human body of science. These four processes together form a powerful endogenously controlled scientific capacity to transform its inherited structures.

Conclusion

Thus we arrive at a (barest) sketch of the richly complex, self-directingly endogenously controlled, deeply anticipatively organised VSR process that is science, the one whose products, both descriptive and procedural, are Campbell’s correlational vicariants. And we have been able to see just how inadequate a purely selectionist account of these cognitive features is, neglecting as it does the higher order organisation of anticipative control. Vicariance, recast as indication, is an organisational control

property which selection theory per se cannot illuminate. It is Campbell's inability to distinguish between high order adaptability and low order adaptation, and between their various organised forms, that is, e.g., behind his inability to distinguish science from religion, sciences from one another and religions from one another. Campbell is joined in this by most other accounts of these processes. In science, e.g., the theory-ladenness of facts, far from being a competing constraint to objectivity, is evidence of the extremely high level of endogenous control depth in science that permits very powerful focusing upon epistemically relevant features for learning (cf. benefits of high order controlled VSR above). A single experiment in physics, e.g., may involve appeal to all or most of physics in its design and data interpretation, including the theory under test (Hooker 1975). In this respect, theory ladenness is an enabling constraint for science which ought to be embraced (with due caution) rather than suppressed. The introduction of a principled account of organisation and control opens up a new vista for understanding both evolution and cognition, in particular science, and their delightfully subtle and powerful natures.

References

- Bickhard, M. H. [1993], "Representational content in humans and machines," *Journal of Experimental and Theoretical Artificial Intelligence*, **5**, 285-333.
- Campbell, D. T. [1974], 'Evolutionary epistemology' in Schilpp, P. A. (ed.) *The philosophy of Karl Popper*. La Salle: Open Court.
- Campbell, D. T. [1997], 'From evolutionary epistemology via selection theory to a sociology of scientific validity', *Evolution and Cognition*, **3**, 5-38.
- Christensen, W.D., Collier, J.D. and Hooker, C.A. [1997], 'Autonomy, adaptiveness and anticipation: towards foundations for life and intelligence in complex, adaptive, self-organising systems', in preparation.
- Christensen, W.D. and Hooker, C.A. [1997], "A Contextual, Embodied, Organisational Theory of Cognition: Autonomous Systems and Self-Directed Heuristic Policies", submitted.
- Churchland, P.S. and Sejnowski, T. [1992], *The computational brain*, Boston: Bradford/MIT.
- Hooker, C. A. [1975], 'On Global Theories', *Philosophy of Science*, **42**, 152-79; reprinted in Hooker, C.A. *A realistic theory of science*, Albany: SUNY Press 1987.
- Hooker, C. A. [1995], *Reason, regulation and realism*, Albany: SUNY Press.

Notes

1. Cf. Bickhard (1993). The important point for the present argument is the introduction of a regulatory context to analyse an epistemic relation. See Christensen/ Collier/Hooker (1997), Christensen/Hooker (1977) and Hooker (1995).
2. Campbell does discuss some contextual issues under the banner of 'co-selection' which we shall discuss below, but this does not address the regulatory context of vicariance.
3. Campbell does refer to the role of mutual interaction in equivocation reduction (section 6), but with respect to VSR epistemology this stands as an ad hoc insight without systematic basis.
4. A fuller account is provided in Christensen/Hooker 1997.