

## Holism and Eclecticism in the Theory of Concepts

Nick Chater

*Department of Psychology, University of Edinburgh, Edinburgh, UK*

Mike Oaksford

*Cognitive Neurocomputation Unit, Department of Psychology,  
University of Wales at Bangor, Bangor, UK*

Howard (1992) defines concepts as the information that a person has about a category, and argues for an eclectic theory of concepts on the basis of this definition. We argue that this definition is unacceptable and hence that eclecticism does not follow. First, the definition is circular as it stands. Secondly, when it is modified to avoid circularity, it implies conceptual holism, according to which concepts are not useful explanatory constructs in psychology. Thirdly, we argue that Howard's argument relies essentially on this unacceptable definition: alternative accounts of concepts, namely categorisational or representational views, do not support it. Having countered the argument for eclecticism, we then argue against it directly on methodological grounds.

### INTRODUCTION

Howard (1992) argues as follows:

1. A concept is everything a person knows about a category. We call this the "concepts as knowledge" account of concepts.
2. Knowledge about categories can be almost arbitrarily varied.
3. Concepts will therefore be equally varied [from (1) and (2)].
4. No single theory can capture all concepts, since they are so varied [according to (3)]; many theories will be required.
5. Current theories of concepts should be seen as complementary rather than contradictory [since (4) means that no single theory is likely to be correct]. We call this theoretical "eclecticism".

---

Requests for reprints should be addressed to Nick Chater, Department of Psychology, University of Edinburgh, Edinburgh, UK.

6. An important research goal is providing a taxonomy of concept types, to elucidate how these are related.

Howard goes on to provide specific taxonomies of concept types, as a framework for future research.

We shall argue that this argument is not persuasive at a number of points. First, our principal aim is to show that (1), the definition of concepts, is not acceptable, and hence that the argument does not get started. As stated, the definition is circular, and if it is modified to avoid circularity, it implies a conceptual holism, which, we argue, rules out concepts as theoretical terms in psychology. In the next section, we argue that the argument for eclecticism relies essentially on this unacceptable definition—alternative views of concepts do not license Howard's argument. Finally, we argue that eclecticism should be strongly resisted on methodological grounds.

### CONCEPTS AS KNOWLEDGE

Despite the enormous amount and variety of work which goes under the banner of "concepts research", there is little agreement concerning what kind of thing a concept is. There are differences of emphasis across research areas. Within cognitive psychology, concepts are often identified with certain kinds of mental structures, variously definitions, prototypes or sets of exemplars (e.g. Medin & Smith, 1984); within developmental psychology, concepts are often viewed as rather abstract aspects of the child's theorising about the world (Carey, 1988); within animal psychology, concepts are usually viewed as primarily concerning the classification of perceptual stimuli (Herrnstein, Loveland, & Cable, 1976); and so on. These boundaries are not hard and fast, however—within each area, many different views are represented, and these and other views are combined and interrelated in many different ways.

Howard cuts through this theoretical tangle by offering a crisp and attractive definition of a concept—as everything that a person knows about a category. This definition is attractive in so far as it is sufficiently broad to be compatible with most other views of concepts and hence offers the possibility of providing a unifying framework for the study of concepts. If anything, it may be a little too broad, allowing a "Pandora's Box" concept, a concept corresponding to the knowledge of the 26 letters of the English alphabet and concepts corresponding to proverbs.

The real problem with the definition is that, as it stands, it is circular. Concepts are defined in terms of knowledge; and this knowledge presupposes the concept being defined. So, for example, Howard identifies the concept TREE with the knowledge that a person has about the category of

trees. But to have any knowledge about trees presupposes that the person has the concept tree (after all, as Howard notes, concepts are the building blocks of knowledge). So, to apply Howard's definition, and pick out the relevant knowledge that a person has about trees, presupposes the existence of the TREE concept.

This circularity is particularly evident if we consider cases in which it is not clear whether or not a particular concept should be ascribed to an agent. So, for example, consider the question of whether a particular child or animal has the concept TREE. According to Howard's definition, to have a concept is to have some knowledge about the corresponding category—in this case, to know something about trees. This puts us no further forward, since if the animal or child knows things about trees, they must necessarily have the concept TREE. So explaining concepts in terms of knowledge has not helped elucidate what it is to have the concept, but simply raises the original question again.

It might seem that a possible reply to this charge of circularity is that knowledge about a category need not presuppose the relevant concept. But this view, if accepted, would undermine the significance of concepts research entirely. Concepts would no longer be the building blocks of cognition, but would have to be assigned some other, more marginal, theoretical role. Certainly, it is clear that Howard would not advocate this defence, since he does see concepts as the building blocks out of which knowledge is constructed.

The problem is that concepts are being asked to play two incompatible theoretical roles: as the representational building blocks of knowledge about a category, and as bodies of knowledge about that category. It is the conflation of these two roles which leads Howard to argue that since a person can represent a prototype, a set of exemplars or a definition of any particular category, that the prototype, exemplar and definitional views of concepts must each be correct. The flexibility that Howard notes is flexibility in the body of knowledge stored about that category. By contrast, in the theory of concepts, the question concerns the nature of the representational building blocks in terms of which this knowledge is stored—whether these building blocks are prototypes, sets of exemplars or definitions.

Although Howard's definition is circular as it stands, the "concepts as knowledge" view can be modified to avoid circularity, by defining a concept not in terms of knowledge of the corresponding category, but in terms of its interaction with the rest of the agent's knowledge. That is, concepts are defined in terms of their role within a person's theory of the world (or some aspect of the world). Circularity is avoided because concepts are defined by their role in an entire knowledge structure, rather than in reference to particular knowledge in which they are involved. This

view is closely related to "theory-based" approaches to concepts (Medin & Wattenmaker, 1987; Murphy & Medin, 1985) and the idea that concepts are characterised by their role in scripts embodying a person's knowledge (Fivush, 1987). However, we shall argue that this modified position falls victim to conceptual holism, which renders "concept" a theoretically useless term for psychological explanation.

According to this position, concepts are defined in terms of a wider body of knowledge or theory; but that knowledge or theory presupposes the concepts in terms of which that knowledge or theory is couched. It seems that within a given theory, each concept presupposes all the other concepts in the theory. This appears to give rise to a circularity of a slightly less direct sort: each concept is defined in terms of other concepts which are, in turn, defined in terms of the original concept. Moreover, given that general knowledge appears to be seamlessly interconnected (e.g. Fodor, 1983) and not decomposable into a number of entirely distinct theories, scripts or schemas, the theory upon which common-sense concepts depend will simply be the whole of general knowledge. This implies that every common-sense concept is defined in terms of every other common-sense concept. Thus, it appears that circularity has not been avoided after all.

There is, however, a familiar trick within philosophy, using so-called Ramsey sentences, for breaking out of this more global kind of circularity. The trick is to accept that one concept cannot be defined within defining all the others, and simply to define the entire theory and all the concepts it presupposes at once. [This trick is used in defining propositional attitudes in terms of each other (Lewis, 1970; 1972; Loar, 1981) and, relatedly, in conceptual role semantics (Block, 1986).]

This modification does not undermine the argument for eclecticism. Step (2) must be revised to state that interaction with an agent's knowledge structures is very varied from one concept to another. Since concepts are now defined in terms of these interactions, this means that step (3), that concepts themselves will be very varied, follows directly, and hence Howard's argument for eclecticism still applies.

But circularity has been avoided only by allowing conceptual holism—concepts cannot be characterised independently, but only as part of an entire knowledge structure or theory of the world. It is important to recognise just how radical are the consequences of holism. It implies not just that a concept, such as TREE, cannot be fully described without taking account of its connections with relevant collateral information; it implies that the concept TREE, like any other concept that a person has, presupposes the concept holder's entire knowledge or theory, including all their other concepts. What this means is that concepts are not building blocks of cognition at all; for the very identity of each concept embodies the nature of the entire cognitive system.

If concepts can only be understood holistically, then concepts research is a futile endeavour. It will only be possible to understand a single concept, such as TREE, by understanding the knowledge structure or theory in which it is embedded (which, as noted above, is likely to be all of general knowledge), and hence all the concepts that this knowledge presupposes. That is, the study of any specific concept is the same as the study of any other specific concept, and is the same as the study of general world knowledge and how it is represented. In particular, the attempt to study individual concepts experimentally is misguided. Furthermore, concepts, construed holistically, cannot be the building blocks of cognitive theory; they can only be ascribed when cognitive theory is complete.

In addition to these general difficulties, a number of more specific problems flow from conceptual holism:

1. There is no account of how concepts can be acquired. Concepts cannot be learnt piecemeal, since possessing any particular concept requires possessing all the rest (and all the collateral knowledge in which they figure, too). For the holist, concept acquisition can only occur, as it were, at a single gulp (see Dummett, 1973; Fodor & Lepore, 1992). Of course, it is open to the concept holist to argue that cognitive development can indeed only be understood *in toto*, that cognitive development is, perhaps, rather like scientific theory change (e.g. Carey, 1988; Karmiloff-Smith, 1988) and that concepts are not acquired piecemeal at all. But if this is so, studying the acquisition of concepts is a contradiction in terms, since, for the holist, the integrity of a concept does not survive alternation in the theory in which it is embedded.

2. The problem of concept acquisition is a special case of the problem that, if concepts are defined holistically, concepts lose their integrity if there is any change to the knowledge structure or theory in which they are embedded. So, for example, learning a new fact about trees not only changes the concept TREE, but will change *all* other associated concepts.

3. Just as a holistic view of concepts rules out the same individual having the same concept over time, it rules out the possibility that two individuals can share the same concept. Thus, since we all know different things about trees, it follows that we all have different concepts of TREE. So there can be no such thing as *the* concept TREE at all; the concept associated with trees will be different for every individual.

These consequences would not be quite so disturbing if there were some way of measuring similarity between concepts in different theories—so that, for example, it would be possible to compare the stock of concepts before and after the new fact was learned, and to match them up, one by one, and see that each had changed only very slightly. Unfortunately,

however, there are no suggestions about how such a similarity measure might be constructed. And, furthermore, where holism has been postulated in other domains, the relation of similarity is generally thought to be ill-defined. So, for example, Quine (1960) assumes that the meaning of terms of natural language is to be defined holistically, and concludes that translation between languages is not determinate; in other words, there is no determinate relation of similarity between the individual words of two languages. Precisely similar arguments could be used to establish an indeterminacy of "conceptual" translation, to show that there is no determinate relation of similarity between the concepts involved in two theories. It is, of course, true that not all philosophers agree that translation is indeterminate—but this is not because they believe that holism is compatible with a well-defined similarity measure, but because they deny the holistic premise. In short, as far as anyone knows, holism cannot be made more palatable by adverting to a notion of similarity between holistically defined entities, because such similarity is not well-defined (see Fodor & Lepore, 1992, for an extensive discussion of why there is no definition of similarity for holistically defined entities).

This means that not only do concepts vary over time and between individuals, but that concepts cannot be compared across times and individuals. For the conceptual holist, concepts are *incommensurable* (to borrow a term from another domain where holism has been suggested, the philosophy of science) across times and people. This means that there can be no psychology of concepts, since studying concepts, in isolation, is simply incoherent.

### ALTERNATIVE ACCOUNTS OF CONCEPTS

We have seen that Howard's "concepts as knowledge" account, which is the cornerstone of his argument for conceptual eclecticism, is circular as it stands, and that a more sophisticated incarnation of his argument falls into the mire of conceptual holism. We have yet to show that Howard's argument depends essentially on this premise, however. In this section, we discuss briefly two alternative accounts of concepts, which we call the representational and categorisational views. We argue that neither of these positions support Howard's argument for eclecticism.

The *categorisational* view is that to possess a concept is to be able to categorise or classify the world in a particular way. So, for example, to have the concept TREE is to be able to distinguish trees from non-trees. This view appears to be implicit within much research on "concept learning", both within comparative psychology (D'Amato & Van Sant, 1988; Herrnstein et al., 1976) and, to some extent at least, in concept acquisition studies in cognitive psychology (Bruner, Goodnow, & Austin, 1956; Shep-

ard, Hovland, & Jenkins, 1961) and the study of the concepts of prelinguistic infants (Goren, Sarty, & Wu, 1975; Johnson & Morton, 1991). (Although, of course, the mere use of discrimination paradigms does not entail that a categorisational view of concepts is presupposed; categorisation performance can, of course, be interpreted as *evidence for* rather than as being *constitutive of* possessing a particular concept.)

This definition is non-circular, since concepts are defined in terms of an *ability*, which can be assessed without presuppositions concerning knowledge or concepts. Furthermore, it avoids holism since each concept is independently associated with a distinct classification ability (although this position does have other difficulties; see Chater & Heyes, submitted).

If concepts are defined in terms of classification, Howard's argument for eclecticism does not apply. A uniform model of classification behaviour is entirely compatible with the diversity of knowledge of categories to which Howard adverts. We do not require a multitude of accounts of concepts to appreciate that "... an individual could have sorts of information about a given category; knowledge of just one exemplar, of three exemplars or a set of defining features or no information at all". The same categorisational mechanism could *establish* the concept in all cases, whatever information may subsequently be stored concerning that concept. Indeed, one of the appeals of the categorisational view is that it allows the possibility that animals, children and adults may have some of the same concepts, when clearly they have very different knowledge.

The *representational* view is that to possess a concept (such as TREE) is to be able to represent the corresponding property (the property of being a tree) in a system of internal representation (Fodor, 1975; 1981). This position is orthodox within cognitive psychology where the fundamental question has been describing the *structure* of these internal representations—whether they are prototypes, definitions or sets of exemplars. According to the representational view, concepts are the building blocks of the internal representations in terms of which knowledge is stored in memory, in the same way as words are the building blocks of the natural languages that we use to store information externally.

*Pace* Howard, the variegated character of knowledge does not imply that the representational system that expresses this knowledge is itself variegated. Natural language is a clear counterexample. It is a uniform representational system, but the information which can be represented in it is arbitrarily diverse. In particular, natural language can represent particular instances of a concept, it can express abstract regularities across many examples, it can express definitions, it can represent prototype structure, and so on. Indeed, what is impressive about language-like compositional representational systems is precisely that they are able to represent an immense diversity of information using a uniform representational basis.

Hence, the diversity of knowledge gives us no reason to assume that the system of internal representation in which that knowledge is couched is any less uniform than natural language. A representational view of concepts will not, therefore, support Howard's argument for eclecticism.

We have argued that if Howard's "concepts as knowledge" view is given up, and a different account of concepts is adopted, the arguments for theoretical eclecticism cannot be sustained. We now argue against eclecticism directly.

### AGAINST ECLECTICISM

In the study of concepts, it is standardly assumed that a single account of all concepts can be provided. Eclecticism recommends that diverse theories should be reconciled, rather than set in opposition. It is in this spirit of reconciliation that Howard's taxonomies are presented: some concepts may be sets of exemplars, some may have Classical definitions, some may have some kind of prototype structure, and so on. Yet this position strongly violates Occam's razor: that the simplest and most general explanations are favoured, and more complex explanations only invoked as a last resort.

In concepts research, adherence to Occam's razor is reflected in the typical theoretical claim that each account of concepts has the widest possible application. It is difficult enough to test the various accounts of concepts empirically even given this interpretation; if many different theories are allowed in different measures and combinations, the theorist will have so many degrees of freedom that empirical considerations may be very difficult to bring to bear at all.

It may of course transpire that some kind of theoretical pluralism is required to spell out the mental processes under study in concept research. There may be no single simple story to tell. Even so, eclecticism will still prove a poor methodological strategy for uncovering this tangle, because it blends together, and attempts to arrange satisfactorily, current theories in a harmonious account [we are arguing that (5) above does not follow from (4)]. Why should a correct pluralist solution happen to consist of a felicitous conjunction of existing views? Since, as noted above, such a compound view will be difficult to refine by empirical research, if eclecticism were adopted, it might be very difficult to move beyond a motley assortment of currently popular theories of concepts.

If, on the other hand, there is vigorous competition between existing positions, those which are inadequate will be rejected or refined, and there will be pressure to develop new theoretical ideas. Later, it may become apparent how some of the results of such competition can be reconciled into a complete account, or a single account may prove sufficient.

It is not easy to think of an example from another scientific domain in which blending competing theories would have proved a successful strategy. One possible exception might seem to be in the debate between the wave and the corpuscular theories of light. However, even here the sophisticated understanding of wave-particle duality would surely never have emerged without the cut and thrust of debate. If it had been agreed that light was partly wave-like and partly particle-like from the outset, debate would have been quenched and progress would have been minimal. What could have been agreed by rival theorists at the beginning of modern physics would simply have been to put their differences to one side and grant that both views had their merits. Only by not settling for such an easy resolution, was it possible to attain the kind of deep theoretical resolution between the two positions which was made possible by quantum mechanics. The moral is that the apparently inconsistent theories of concepts should *remain* rivals until some genuine theoretical resolution can be found or one or other theory gains the upper hand.

Manuscript received September 1991  
Revised manuscript received October 1992

### REFERENCES

- Block, N. (1986). An advertisement for a semantics for psychology. In P. French, T. Uehling, & H. Wettstein (Eds), *Midwest studies in philosophy. Vol. 10: Studies of the philosophy of mind*, pp. 615-678. Minneapolis: University of Minnesota Press.
- Bruner, J.S., Goodnow, J.J., & Austin, G.A. (1956). *A study of thinking*. New York: John Wiley.
- Carey, S. (1988). Conceptual differences between children and adults. *Mind and Language*, 3, 167-181.
- Chater, N. & Heyes, C. (submitted). *Animal concepts: Content and discontent*.
- D'Amato, M.R. & Van Sant, P. (1988) The person concept in monkeys. *Journal of Experimental Psychology: Animal Behavior Processes*, 11, 35-51.
- Dummett, M. (1973). *Frege: Philosophy of language*. New York: Harper and Row.
- Fivush, R. (1987). Scripts and categories: Interrelationships in development. In U. Neisser (Ed.), *Concepts and conceptual development*, pp. 234-254. Cambridge: Cambridge University Press.
- Fodor, J.A. (1975). *The language of thought*. New York: Crowell.
- Fodor, J.A. (1981). The current status of the innateness controversy. In *Representations*. Brighton: Harvester Press.
- Fodor, J.A. (1983). *The modularity of mind*. Cambridge, MA: Bradford Books/MIT Press.
- Fodor, J.A. & Lepore E. (1992). *Holism: A shopper's guide*. Oxford: Basil Blackwell.
- Goren, C.C., Sarty, M., & Wu, P.Y.K. (1975). Visual following and pattern discrimination of face-like stimuli by newborn infants. *Pediatrics*, 56, 544-549.
- Herrnstein, R.J., Loveland, D.H., & Cable, C. (1976). Complex visual concepts in the pigeon. *Science*, 146, 549-551.
- Howard, R.W. (1992). Classifying types of concept and conceptual structure: Some taxonomies. *European Journal of Cognitive Psychology*, 4, 81-111.

- Johnson, M.H. & Morton, J. (1991). *Biology and cognitive development: The case of face recognition*. Oxford: Basil Blackwell.
- Karmiloff-Smith, A. (1988). The child is a theoretician, not an inductivist. *Mind and Language*, 3, 183-196.
- Lewis, D. (1970). How to define theoretical terms. *Journal of Philosophy*, 67, 427-446.
- Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*, 50, 249-258.
- Loar, B. (1981). *Mind and meaning*. Cambridge: Cambridge University Press.
- Medin, D.L. & Smith, E.E. (1984). Concepts and concept formation. *Annual Review of Psychology*, 35, 113-138.
- Medin, D. & Wattenmaker, W. (1987). Category cohesiveness, theories and cognitive archeology. In U. Neisser (Ed.), *Concepts and conceptual development*, pp. 25-62. Cambridge: Cambridge University Press.
- Murphy, G.L. & Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
- Quine, W.V.O. (1960) *Word and object*. Cambridge, MA: Harvard University Press.
- Shepard, R.N., Hovland, C.I., & Jenkins, H.M. (1961). Learning and memorization of classifications. *Psychological Monographs*, 517(75), 1-42.