

CHAPTER SIX

CONCLUSION

6.00 CHAPTER OVERVIEW

The aim of this chapter is to present answers to the two questions posed by the thesis, and from them draw some conclusion about how people organize and comprehend their world. Section 6.01 summarizes the main findings of the thesis.

Section 6.02 describes the main findings of the thesis with regard to the first question (why do concepts arise), and describes the main conclusions that follow. The section discusses the results from the empirical studies of the thesis with respect to whether or not our concepts arise as a result of the already-existing categories in the environment. The possibility that our conceptual representations reflect the metaphysical structure of the environment was not supported by the results of Experiment 1, and the possibility that they arise from its physical structure was only partially supported by the results of Experiment 2. These experiments eliminated the structure of the environment (be it metaphysical or physical) as the sole basis for concepts. However, in Experiment 2, the small size of the semi-partial correlations between predictors and criterion for representation, would indicate that a third factor was at work during categorization, and the possibility that it was theoretical knowledge was tested in Experiment 3.

Section 6.03 describes the main findings of the thesis with regard to the second question (why do we have the categories we do, and not others). It is concluded that we have the categories we have because our theories of the world dictate what *appears* similar to us. In short, we have certain categories (and not others) as a result of the concepts we construct about the objects and creatures in our world. Based on the results of Experiment 3, it is concluded that a person's concept of a category arises through his or her beliefs and theories which explain the outside world, such as the creatures and objects in it. More specifically, it is pointed out that, as the stories based on the two-tier representation model showed, people can draw inferences about an object's or creature's category identity on the basis of their background knowledge of functions, needs, and appearance.

Section 6.04 argues that the general purpose of concepts is to intentionally understand the world around us, using concepts to build up our knowledge. More specifically, by incorporating the properties of stability, coherence and flexibility into their theories, explanation-based views of concepts and categories can provide an account of a concept's functions. This section looks at some of the unexpected findings of the empirical studies, and attempts to draw conclusions about them by using an explanatory-based account of concepts. The main conclusion is that, in order to attain some degree of plausibility, concepts need to include some subjective knowledge (as well as normative) into their content.

The section also looks at some of the drawbacks (subjectivism and relativism) inherent in a "constructivist" theory of concepts which sees their main function as one of serving to actively construct reality. Theoretical resolutions suggested by Goodman (1984) are discussed.

6.01 SUMMARY OF THE MAIN FINDINGS

The results of Experiment 1 showed that all three category-types had graded structure, with differing degrees of gradience across their levels of production frequency. The superordinate types had the steeper gradience in that more of their ten levels were significantly different from one another, than in the other two types.

The results of Experiment 2 did not support the four assumptions made by the Unitary View's approach to categories, their representation, structure and process. In the study, participants were shown to have dual representation of the categories, storing more than one kind of information about them. Their categories had clear-cut boundaries, and their items were not categorized in a single, one-stage decision process. Overall, these results support a Binary View of categories, which holds that categories are the result of concepts and theories we have about the world.

This possibility was tested in Experiment 3, which showed that participants found the most relevant information about animals to be concerned with goals, needs and preferences, whilst about artifacts, the information concerned their functions and abilities. Contrary to prediction, the study found

no differences between the importance of physical alterations (condition 2), and explanations for those physical alterations (condition 6). Since the claim is that the conceptual core consists of theories and beliefs about a category, then condition 6 should have achieved a greater rate of change in people's judgments (as compared to the control condition), than what was achieved in condition 2. However, the prediction that explanatory-based conditions would elicit differences in performance between the judgment tasks was upheld. Post hoc tests did indicate that, in condition 6, participants were using physical information for similarity judgments and explanatory information for categorization judgments. This was not the case in condition 2, where solely physical appearance was described, and people showed no differences in performance among the three judgment tasks.

6.02 HOW DO CONCEPTS ARISE?

The first question posed in this thesis was, how do concepts arise? The main point derived from the results of Experiment 1 was that people can share the same mental representations of what members might constitute a category, whether natural, property, or ad hoc category-types are involved.

Experiment 1 showed that people's generation of exemplars is not driven by strictly mathematical rules, so the possibility of a logically-based external environment which determines our categories was eliminated in this study. Consequently, it is concluded that the nature of the representation of the category-members does not involve a mathematical structure. If it did, then the distribution for the number of different items produced as exemplars of the category would have been a function of the increasing or decreasing number of people doing the producing. Instead, subjects produced a distribution which seemed to be based upon typicality of the items, with the number of different items produced lessening as the number of people increased. This suggests that the majority of people agreed upon which were the most typical exemplars of a category. Most people, when requested to produce the best exemplars of the category first, produced a small range of the same exemplars first. This was strong evidence for a stable membership structure based upon typicality of the instances, and for the argument that subjects were consulting a common or shared category representation, in a consistent and predictable fashion.

In Experiment 1, the individual instances were not organized in a haphazard manner, which would be the case if participants were using an algorithmic formula to generate them, so categories cannot be represented in an abstract fashion. The results imply that Anderson's (1990) theoretical assumption about a world structured by Bayesian formulae was mistaken, and so the source of meaning does not lie in such a world. For a more detailed description of this issue, see Appendix J.

If concepts originated mainly in the physical structure of the external environment, then the knowledge they contain should involve only normative knowledge about how the world (its objects and creatures) is "carved up". Experiment 2 showed decisively that this is not the case. The experimental results suggest that there was a third factor at work, one which was not tested for in Experiment 2, and that people were consulting this factor for their membership decisions and their ranking tasks. Questions are raised by the small size of the semi-partial correlations between predictor task (for example, measures of typicality or frequency of direct experience) and the criterion for mental representation of the category. These suggest that both typicality and direct experiential information constitute only a small amount of the knowledge represented about a category, leaving a great deal of the category representation unexplained. It is argued that this third factor is conceptual knowledge, such as that described by Landau (1982), which consists of knowing the general facts or common usage of what being a grandmother means, for example, as well as the specific biological definition.

A view of concepts which makes some allowance for people's theories, beliefs, social knowledge about the category takes the nature of concepts one step further than a strictly environment-based knowledge. Such a view proposes (as Kant did, see Appendix B) that sensory experience or conceptual understanding alone are not sufficient: both are required in concept representation. Kant (1787) suggested a mediation of the sensibility and the understanding by imagination. Such a view also allows for a much more active role to be played by the categorizer. It is pointed out that, as the stories of Experiment 3 based on the two-tier representation model showed (condition 5), people can draw inferences about an object or creature's category identity on the basis of their background knowledge of functions, needs, and preferences.

Medin (1989) has suggested that some of the stability evident in theory-based concepts might derive from certain tendencies people show in their reasoning and problem-solving behaviours, which are evident across cultures. This is consistent with Kant's (1787) argument that we all have certain knowledge, a priori, by virtue of having human minds, and that such a priori knowledge precedes all reasoning. Kant, in turn, based his ideas on Hume's discovery that certain relations among things in the real world cannot be attributed to events, but rather were "mental constructions" projected onto an "objective world" (Bruner, 1986).

One conclusion drawn from the empirical studies of the thesis is that concepts are constructed, not derived from metaphysical categories which we discover to be already existing in the universe, or developed from similarity structures inherent in the environment. Some of the stability in concepts is derived from people's innate tendency to perceive certain correlations between features in the environment, and then place their own "mental constructions" upon them. What those mental constructions might be, however, depends not only upon the cultural norms of the categorizer (Smith, 1984), but also on his or her individual experiences. In short, upon the accumulated knowledge of the "given" world in which the categorizer lives.

6.03 WHY DO WE HAVE THE CATEGORIES WE DO?

The second question investigated by the thesis, and discussed in chapter 2, asked why do we have the categories we do, and not others. One possible answer is the notion that the function of categories is for cognitive economy in the prototypical representation of the outside world. This view is known as the unitary approach, because it assumes that categories reflect something *direct* about the structures or essences of the outside world. The alternative answer is that we have the categories we do, as a result of the diagnostic criteria, which is contained in our concepts, and which we impose upon the outside world.

Rosch (1978) has argued cogently that we have categories in order to achieve cognitive economy in our mental representations of similarity structures inherent in the outside world. The cognitive economy principle claims that it is desirable to minimize effort, and the way to do this is by people being able to treat a grouping of objects or creatures in the same way. Furthermore, maximum

storage of information is achieved by the use of categories, so that our memories do not become swamped with irrelevant details about the individual object or creature (Murphy, 1993b). Cognitive economy is a clearly credible purpose for the existence of categories, but like prediction of features (the purpose for categories advocated by Anderson, 1991a), it also is a limited view of the variety of functions open to categories. People use categories for other purposes also, not only that of organizing the creatures and objects of the physical world they live in.

Experiment 2 tested the assumptions contained in a unitary approach to categories, which assumes similarity to be the basis for categorization. The results did not support the unitary approach, but supported the dual representation (or binary) approach to category representation, which has the drawback of not being cognitively economical in Roschean terms. Where people's processes were concerned, the unitary approach claims that categorization consists of a single stage of computation of similar, characteristic features. Experiment 2 showed this not to be the case, with participants undertaking a two-stage processing of items, as evidenced in their response times.

Furthermore, Experiment 2 showed that people do not treat all groupings of objects/creatures in the same way, nor did the different category-types (superordinate, property, ad hoc) consist of the same information. Each had single and dual representations. Prototypicality was present in all three normatively-based category-types, but it was not the sole basis for category representation, thus indicating that cognitive economy of information storage (in the Roschean sense) was not really being achieved.

In Experiment 2, where representation of the natural superordinate types was concerned, "cognitive economy in the representation of the physical environment" seems to be an appropriate function of categorization. The importance of personal experience and frequency of category-context in our understanding of natural categories was evident in superordinates, by the fact that the frequency-of-instantiation task was found to be a significant predictor of representation also. This finding suggests that the unitary approach's view of natural categories is too simple, and that physical appearance is not the only information encoded about them.

It must be concluded that Rosch's (1978) answer to the question is inadequate. We do not have the categories we do for the sole purpose of representing physical information in a cognitively economic fashion. This is most evident in ad hoc categories, which Experiment 2 showed were represented most uneconomically by three different kinds of information (frequency of experience, typicality, and membership criteria reflected in the Ranks variable).

The coherence of ad hoc types cannot be explained according to a unitary approach to categories. In the results from Experiments 1 and 2, the ad hoc types showed a significant gradient in structure which, though slight, indicated that they are categories in their own right. As such, they raise the question of how they can have a graded structure which is *not* based upon similarity of their items. Barsalou (1983) suggested that each exemplar signified how well it fulfilled the goal of forming the category in the first place. For example, in the category *Things to save from a burning home* the goal would be to list items of value. Thus, the degree of value inherent in the item (whether sentimental, monetary, or an intrinsic value such as life) would provide gradient to the category's membership structure.

The presence of content information (as well as structural typicality) in the representations of the three category-types would suggest that, in all three types, people were using a knowledge-based concept to generate their exemplars. It is argued that the coalescing of the various exemplars of any of the three category-types is due to the membership criteria specified in the concept for the category. This criterion may have nothing to do with the physical similarity of the items, but a great deal to do with beliefs about why those exemplars *appear* similar.

Concepts ensure coherence of structure by "diagnosing" objects and creatures, on the basis of something other than their physical appearance (Rips, 1975; 1989; Smith, 1989). A view of concepts as consisting of diagnostic beliefs about a domain would provide an explanation for the property of coherent structure in categories consisting of physically dissimilar objects (as in ad hoc categories). Perception of similarity is not fixed, but contingent upon the theory-based concepts or beliefs which people hold about a group of objects as a category. Thus, a woolly lamb and a woolly dog might look physically similar, but one meets the specification of what is relevant or important for belonging to the "doggy" category, and the other does not. Consequently, another reason or

role for concepts is to diagnose input from the external environment, in a way that category membership decisions are based on more than physical appearance or rules.

In Experiment 3, the stories in condition 6 were based on Medin's (1989) descriptions of contagion thinking and homeopathic problem-solving (see chapter 1). This story condition elicited the second highest rates of changed judgments (as compared to the control condition), out of the six experimental conditions. The subjects had no difficulty at all in comprehending the stories, or in drawing certain inferences. For example, one animal story (story item 3, see Appendix I) describes a change from lizard to owl. When a lizard stands too close to a bird on a rock, it catches the bird's contagious microbes, which infect its cold blood so that the lizard is now warm-blooded. Whilst the credibility of the story is low, subjects still found it comprehensible, and the information contained in it important enough to effect a significant change in their judgments, so that they considered the lizard to have changed its membership into the category of owls.

The results of Experiment 3 indicated that the explanatory concepts we construct about the objects and creatures in our world are the decisive factor in cohesiveness. However, they also showed that appearance is an important component of that diagnostic process, and such appearance need not necessarily be based upon physical similarity. In short, similarity is dynamic and sensitive to context-effects.

Finally, on the issue of whether concepts or categories come first (Hampton & Dubois, 1993), the thesis must take the position that it is concepts which determine categorization rules and decisions. We do not have the categories we do because they were already inherent in the environment, and were there waiting to be discovered. We do not have the categories we do in order to predict features or represent the physical environment in a cognitively economical fashion. We have categories because of the theories we hold about the physical world around us, using such beliefs and theories to *diagnose* creatures or objects and assign them to some category or other. We have the categories we do because we are the sort of organisms we are (Wattenmaker, Nakamura and Medin, 1988).

6.04 THE ROLE OF CONCEPTS IN COGNITION

Murphy (1993b) suggests that people use their theory-embedded concepts to interact with their sensory environment and to make sense of it. Thus, a concept model should portray representations of physical appearance as interacting in predictable ways with representations of concepts. Recent research has shifted to investigation of the categorizer, rather than hypotheses about the essential nature of the outside world, or its metaphysical structure. The nature of the world is not wholly ignored, but rather, the focus is upon the interaction of intelligent organisms with their environment (Murphy, 1993b; Rips, 1986; 1990).

The thesis argues that the overall purpose of concepts is to construct some meaning for the world we live in. The three functions of concepts detailed by the theory-based view of concepts are specific, and it is in carrying out these functions that three properties of people's cognitive behaviour become evident. Of the three approaches to concepts examined in this thesis, the view of concepts as theory-embedded is the only one which can explain stability of representation, flexibility of categorization, and coherence of structure. So taking into consideration this account's explanatory capacity, it would seem to be the most valid one of the three. Also, the account should be able to explain some of the unexpected findings from the empirical studies.

6.04.1 A flexible categorization of creatures and objects

Experiment 3 showed that judgments of an animal or artifact could be influenced by the story-context, as could the participants' perception of the animal or artifact. In other words, a person's judgment was not carried out on its own merits, and similarity was not judged strictly on information about physical features.

Experiment 3 compared subject performance in six experimental conditions, asking whether an item should be conceptualized or identified with one concept or with another. The results indicated that similarity relations between an object and its category are not fixed solely by physical information concerning their appearance. For example, in the stories, subjects changed their judgments of an object's similarity to a particular concept when function (condition 4) or essence (condition 3) were varied. Where categorization

judgments were concerned, the object's category identity changed, even when the information consisted solely of details about physical appearance.

Schank, Collins, and Hunter (1986) criticize category induction theories because these assume that the categorizer has access automatically to all relevant and important features. One of their excellent examples includes that of predicting stripes if the animal is a zebra. They discuss some of the problems which might ensue if some zebras do not have stripes, in which case a categorizer might mistakenly take them to be horses. In short, they conclude that, "the importance of a feature is not static, and can only be determined by pragmatic judgments about the features, their context, and the system's goals" (Schank, Collins, & Hunter, 1986, p. 642). This was borne out by Experiment 3, where similarity judgments which included the same physical feature differed according to the story context where that physical feature was described. In short, prototypicality did not seem to be a *fixed* perception in the stories of Experiment 3, either, where variations to an object's or creature's functions or needs influenced subjects' perceptions of its similarity and its category identity. This means that the number of categories an artifact or animal can belong to are more than one, two, or even three. In fact, everything can share at least one feature with something else, and a new category can be formed on the basis of that shared similarity.

The main drawback of the theory-based view of concepts is that the categorization processes which they describe are so flexible as to be almost unconstrained (Barsalou & Medin, 1986). This is because features that are correlated in people's mental representations do not always reflect actual empirical relations in the world, but may derive instead from people's theories or beliefs about the observed relations between the features. Relations between features do not need to be actually observed. They may emerge as a consequence of someone's theories about the world. When a correlation is perceived to exist on the basis of one's theories, but has no basis in empirical fact, it is called an *illusory correlation* (Murphy & Medin, 1985). Since people are said to perceive similarities according to beliefs and theories, anything might be perceived as similar to anything else. So why do we have the categories we do, *but not others*? Roschian theory provides no constraints upon these illusory correlations, since it makes no attempt to explain why certain features (those which went into the mental calculations of similarity) might be chosen to be compared or perceived as being similar in the first place.

The flexibility of similarity relations needs to be constrained, and concepts achieve that constraint by specifying what is relevant *to the category*. Constraints on such flexibility are one of the functions of concepts, insofar that a concept is said to specify what information is *relevant* or *important* to the category (Wattenmaker, Nakamura, & Medin, 1989). In this way, it is claimed, the categorizer makes his or her decision about an instance's membership, by knowing which similar features should be taken into consideration. The question then becomes one of what is the relevant information. Theory-based views of concepts claim that beliefs and theories specify what are the relevant features to take into account during a membership decision. This description of what is *relevant* information does not agree completely with the empirical results from Experiment 3.

6.04.2 A coherent structure for categories

Just any explanatory theory will not do. In Experiment 3, one of the unexpected findings was that story conditions 2 and 6 were equal in their influence upon people's judgments. Just a plain description of an alteration to physical appearance (condition 2), had as much force and influence on judgment tasks as when an explanation was provided for the transformed appearance (condition 6). As the lack of difference between conditions 2 and 6 showed, just any explanation will not do.

The beliefs drawn upon in condition 6, might influence *constructed* similarity, since the similarity judgments in this condition were significantly higher than the categorization judgments. However, this constructed similarity was still not enough to influence categorization *to a greater degree* than the physical similarity described in condition 2. The degree of changed judgments in condition 6 were no different from the degree of change effected by the stories of Condition 2. The explanatory story condition which did achieve this was that of condition 5. The essential difference between condition 6 and condition 5 (which both provided explanations for any transformations) was that the latter provided personal details (about animals) or functional details about artifacts and the needs of their owners, and based its explanation on those.

Of the six experimental conditions tested in Experiment 3, the most relevant information for participants was that contained in the stories of

condition 5. These were stories concerned with an animal's needs, its intentions and preferences, or an artifact's functions and abilities. It is argued that such knowledge, for the categorizer, is embedded in self-reference. The categorizer is using the knowledge gained from his/her own experience of needs and preferences to judge the animal stories in condition 5, (and his/her own experience of artifact functions), to judge the *plausibility* and *credibility* of the stories themselves.

This is not the same kind of relevant information portrayed by the stories in condition 6 (psychological essentialism), though these did gain significant differences from the control condition also. The knowledge contained in the condition 6 stories involved the use of theoretical information which might be gained from a book or from school. It certainly did not involve the use of knowledge which is credible because it is gained directly from being human, and sharing the world with other humans.

Medin (1989) speaks of the need for similarity, to provide constraint upon relations within a category. Goodman (1977) points out that each similarity judgment needs a frame of reference to stabilize comparisons, and avoid the charge of relativism. Perhaps a need for plausibility also exists, to provide a guide to the credibility of a theory, some way by which subjects can judge how believable a theory might be. Plausibility is not the same as coherence; it does not serve the same functions as coherence, that is, diagnosis of membership. A theory might be coherent, but not plausible, as Johnson-Laird (1983) has observed. It is claimed here that plausibility serves to help the categorizer decide whether to accept or reject any explanatory theory, on the basis of whether that theory "fits in" with the version of the world which he or she has created for themselves. By providing a frame of reference, concepts which are embedded in self-reference are fulfilling the function of specifying what is *plausible* and important information to take into account during categorization.

This kind of self-reference (knowledge gained from direct experience) might be argued against on the basis of subjectivism, where everything is relative to everything else. Self-reference might be argued to be too subjective, in which case how can we communicate with one another, if the meanings we construct are completely idiosyncratic? The next section on stability of conceptual representation deals with the need for both subjective *and* normative knowledge.

6.04.3 A stable representation for concepts

The empirical results gained from Experiment 2 suggested that experiential knowledge encoded subjectively is just as important as normative knowledge gained from social norms. At least, the data gained from the use of idiosyncratic stimuli was as useful as that gained from using normative stimuli, and sometimes more profitable in terms of producing a greater number of significant results.

In Experiment 2's regression analyses of category representation, each subject's data were analysed individually. The accountability ratios elicited by each individual's performance when using idiosyncratic stimuli were shown to be equally as high as those elicited by the use of normative stimuli. If the argument were supported that it is the subjective components of information which produce instability, then the accountability ratios of the predictor tasks using idiosyncratic stimuli should have been far lower than when the same tasks were performed with normative stimuli. Instead, the values were roughly equal, irrespective of whether the stimulus words used for a task were idiosyncratic or normative.

In the experiment's "judges' agreement" analyses of category membership, the results showed opposing outcomes in the superordinate and property category-types, depending upon whether the data from the normative stimuli were being used, or those from the idiosyncratic stimuli. Judges agreed about category boundaries when the normative stimuli were used, suggesting that both superordinate and property types had clear-cut and precise membership; but the opposite occurred when idiosyncratic stimuli were used. Also, in the experiment's second analysis of the membership decision response times, the analysis of the idiosyncratic-based data produced a greater number of significant results than the analysis of the normative-based data.

These results of Experiment 2 would suggest that idiosyncratic knowledge gained by direct experience is at least as important as normative knowledge gained from society's norms. Overall, the main conclusion reached from this section of the results is that future studies on concepts and categories might find it profitable (in terms of significance of results) to use not only the normative organization of category members which has been extracted across hundreds of

participants, but also each participant's idiosyncratic organization of his or her category members. Certainly, the same task (but with different stimuli) produced data which often produced differing results.

As was described in chapter 1, Rey (1983) suggested that two kinds of conceptual stability need to be accounted for, in any adequate theory of concepts: within-person stability and between-person stability. Within-person stability is what occurs when a person has the same concept at different times, and it is this within-person stability which provides the basis for the person's conceptual competence (Rey, 1983). For example, a person's understanding of how a sewing machine functions is usually learned through direct experience, often through trial-and-error. It is necessary for the same person to have this same concept at different times. Although the person might be sewing different garments on different occasions, the within-person stability of an understanding of how a sewing machine functions is necessary on every occasion. Otherwise, it would be necessary to re-learn the same mechanisms each time she or he wants to sew something. Within-person stability of concepts, then, is grounded in personal and direct experience of the world.

Idiosyncratic knowledge, beliefs derived from personal direct experience of the world and its creatures, seems to be a valid basis for within-person stability of concepts. Knowledge of something learned at first hand, through direct experience or in an emotion-laden context, is likely to be better recalled or better recognized (than statements of cold fact). This knowledge derived from direct experience or emotions contributes towards stability within-persons, through self-reference. Idiosyncratic knowledge should hold meaning for an individual through its association with one's self.

Between-person stability occurs when different people can share similar concepts at the same and different times, providing the basis for comparisons of norms across people in a society. Normative knowledge of something learnt through indirect, vicarious experience (for example, from a book or at school) may not have the same salience or force as idiosyncratic knowledge. Comparison of participants' performances with normative stimuli, as compared with idiosyncratic stimuli, would suggest that participants may have subjective knowledge associated with their categories, as well as a normative organization.

Theories as a basis for conceptual stability are said to serve as an "anchoring" function for concepts (Medin & Ortony, 1989; Keil, 1989), yet they have been criticized because of their lack of specificity. As Keil's (1989) notion of causal homeostasis implies, concepts and the background theories from which they derive are highly interrelated, so how are they differentiated from each other? The proposal that ordinary, everyday concepts might be represented as beliefs and embedded in theories needs to be more explicit, to avoid the danger of circularity of explanation (Komatsu, 1992). The thesis suggests that, by allowing a concept to have both subjective and normative components, Rey's (1983) conditions for the stability of representation would be achieved. Furthermore, the concept would be differentiated from the background theory by virtue of its subjective component of knowledge, by which the categorizer can judge the plausibility of the theory.

6.04.4 Constructed worlds

Goodman (1984) sees the greatest problem with the idea of "constructed" concepts as being one of recursiveness, the explanatory "buck" keeps being passed *ad infinitum*. He proposed that certain theories (or prior worlds as he describes them) be taken as "given" for the individual, so that the categorizer is constrained by the nature of the world version with which he or she began remaking concepts. To illustrate how this might occur, there is an instructive exchange between Marco Polo and Kublai Khan in Italo Calvino's *Invisible Cities*. It begins when Marco says:

"Sire, now I have told you about all the cities I know."

"There is still one of which you never speak... Venice," the Khan said. Marco smiled. "What else do you believe I have been talking to you about? ... Every time I describe a city I am saying something about Venice."

"When I ask you about other cities I want to hear about them. And about Venice, when I ask you about Venice."

"To distinguish the other cities' qualities, I must speak of a first city that remains implicit. For me it is Venice." (Calvino, 1972, p.86).

By taking certain constructed worlds as "given", whether they be through the categorizer's direct experience or through living in a certain culture, the charge of relativity is avoided. Jerome Bruner (1986) has argued for a constructivist view: that we cannot know an aboriginal reality; that there is none, that any reality we create is based on a transmutation of some prior "reality" that we have taken as given. We construct many realities, and do so

from differing intentions (Bruner, 1986, p.158). We construct our worlds (or realities) out of the myriad forms in which we structure experience, whether it be the experience of the senses, the deeply symbolically encoded experience we gain through interacting with our social world, or the vicarious experience we achieve in the act of reading.

If the general overall purpose of concepts is not only to build, but also create conceptual knowledge, then by constructing concepts we are constructing worlds. Such functions cannot be reduced to simple feature prediction, or the efficient storage and processing of the physical structures of the external world (Barsalou, 1992; Smith, 1989). If there are meanings "incarnate" in the world, Goodman (1984) argues, we transform them in the act of accepting them into our transformed world, and that transformed world then becomes the world with which others start, or that we then offer.