

A THEORY
OF
ABSTRACTION

David Kelley

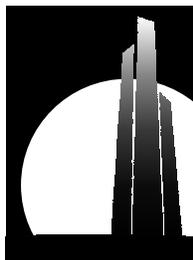
The ATLAS SOCIETY

“A Theory of Abstraction” was originally published in *Cognition and Brain Theory*, 1984, 7 (3 & 4), pp. 329–357.

Copyright © 1984 by David Kelley

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, mechanical, electronic, photocopying, recording or otherwise without prior written permission of the publisher. Information or permission requests should be addressed to The ATLAS SOCIETY, 1001 Connecticut Ave. NW, Suite 830, Washington, DC 20036.

Printed in the United States of America. Reprinted 2010.



The ATLAS SOCIETY

1001 Connecticut Ave. NW, Suite 830
Washington, DC 20036
202-296-7263 (AYN-RAND) (Voice)
202-296-0771 (Fax)
tas@atlassociety.org
www.atlassociety.org

A Theory of Abstraction

David Kelley

Abstract

The model of concept-formation defended here, on philosophical and psychological grounds, is based on the work of Rand (1979). It is abstractionist in the sense that the process of forming a concept derives from the perception of similarities among objects. The process involves two stages, each characterized by a specific mode of attention to perceived similarities and differences. In addition to the local arguments in support of each stage, the general constraints that should be placed on any theory of concept-formation are also discussed.

A Theory of Abstraction

If empiricism is the doctrine that all knowledge of the world derives from the senses, then it would be natural to expect an empiricist to show how concepts can be acquired from perceptual experience. We should expect a good empiricist to have a theory of concept-formation. This expectation is satisfied by the historical empiricists. Both in the Aristotelian tradition, and in the modern tradition from Locke to Mill, one will find many attempts to provide a theory of abstraction—using the term “abstraction” in the broadest sense. Berkeley and Hume, of course, attacked the very notion of “abstract ideas,” but even they tried to explain how words, at least, could come to be associated with classes of particulars and thus function abstractly. In this very broad sense, empiricism is committed by its nature to abstraction, precisely because it rejects the claim that abstract words or concepts are innate.

It is curious, then, that philosophers in this century, despite their widespread commitment to empiricism, have rarely dealt with the issue of concept-formation. The problem of universals has endured as a topic in metaphysics. But questions about how we form concepts to identify universals—questions that were once considered the epistemological dimension of the metaphysical problem—have *not* been a major topic in twentieth-century epistemology. The reasons for this are complex, bound up as they are with the linguistic turn in philosophy, and I cannot do justice to them here. But I do want to comment on two points that I think have been particularly important, in order to clear a space for my primary goal: to present a new theory of abstraction.

4

First, we should recognize the currency of a certain argument against the very possibility of a theory of abstraction. Any such theory holds that the acquisition of a concept depends in some way on the observation of similarities among perceived objects. But then, it seems, we must already possess the concept 'similarity' in order to engage in the process, and so abstraction cannot explain the acquisition of *that* concept. (This is the epistemological parallel of an old metaphysical argument for realism concerning universals: the argument that even if two white objects have no abstract property literally in common, but are only similar, that very fact commits us to the existence of the universal *similarity* (Russell, 1959).)

Consider the case that Sellars (1963) makes against foundationalism. Empiricists have typically held that knowledge of the world rests on a foundation of basic judgments, justified directly by perceptual experience. Sellars' argument against the possibility of basic judgments turns in part on his coherence theory of concepts. The predicate employed in any alleged basic judgment, he argues, has the content it does because of its place in a network of other predicates, and in a network of assumptions, beliefs, and theories in which such predicates are embedded. Thus one cannot know one thing without knowing many (see also Churchland, 1979, Chap. 2). Why can't the concepts employed in basic judgments be acquired directly from perception, through the observation of similarity? Because that would be a circular account:

...if the formation of the association [between words and classes of resembling particulars] involves not only the occurrence of resembling particulars, but also the occurrence of the awareness *that they are resembling particulars*, then the givenness of determinate kinds or repeatables, say crimson, is merely being replaced by the givenness of *facts* of the form *x resembles y*, and we are back with an unacquired ability to be aware of repeatables, in this case the repeatable *resemblance*. [p. 160]

Another example, from a very different context, is Fodor's argument against abstraction (in 1975). As I have argued elsewhere (Kelley and Krueger, 1984), current psychological theories of concepts all presuppose the subject's ability to identify objects as instances of abstract properties. These theories differ on issues of how abstract properties combine to form the representational contents of concepts, but leave unexplained the ability to detect abstract properties themselves. Thus there is nothing in the psychological literature that could provide us with a non-circular ex-

planation of concept acquisition. In effect, the theories commit us to the view that at least certain elementary concepts are innate. Fodor's argument is that some form of nativism is indeed inevitable. For even a theory of abstraction

...will mitigate the nativistic assumptions about *concepts* at the price of nativistic assumptions about similarity. (One cannot use *C is the concept of things sufficiently similar to E_i* [an exemplar] unless one is already in a position to employ *is sufficiently similar to E_i* .) [p. 97]

where the italicized expressions are abstract predicates.

Now it is clear in a general way how the empiricist should answer this argument: he should point out that the process of abstraction does *not* require the conceptual awareness of similarity. It involves the perception of the similarities among perceived objects, but does not require the possession or use of the concept 'similarity.' More generally, a theory of abstraction is intended as a description of a sequence of cognitive states and processes occurring at the preconceptual level. And this brings us to the second point. Epistemologists have shied away from the idea of preconceptual cognition. This animus is sometimes defended in terms of a distinction between causal and justificatory issues in knowledge. Fearing the charge of psychologism, many have held that epistemology should be concerned solely with the justification of beliefs, and that justification is a matter solely of the logical relations among their propositional contents; questions about the cognitive equipment employed in propositional knowledge, particularly about the manner in which it is acquired, are matters exclusively for psychology.

But I think the animus has deeper roots—in the Kantian thesis that percepts without concepts are blind, i.e., non-cognitive. Many philosophers seem to hold that a percept may have an object in a causal sense, because of its causal relation to a distal stimulus, but it cannot have a genuinely *cognitive* content, an intentional content, unless it involves some conceptual recognition of the object, some element of predication, some propositional form. A perceptual *judgment*, in other words, may be a cognitive phenomenon, but a pre-judgmental, preconceptual state is not. (The first part of Sellars' essay is a kind of extended exposition of this assumption.) Again, the psychological paradigm that Fodor defends holds that all cognitive processes are computational, involving some sort of logical relations among the contents of cognitive states. Thus Fodor (1981) claims that for there to be a *cognitive* process of concept acquisition, there must be "a rational relation between the concept acquired and the experience

that occasions its availability” [p. 275]; and he makes clear that by “a rational relation” he means a logical one, a relation which holds between propositional, and thus necessarily abstract, contents.

The widespread appeal of the circularity argument, then, and the disrepute into which theories of abstraction have fallen, is really the result of a package: the premise that all cognition is linguistic or at least propositional, and conversely that there are no prepropositional cognitive states or processes; and the consequent insistence on a sharp distinction between causal and justificatory dimensions of knowledge. I have argued elsewhere (Kelley, 1986) against every element in this package. In this context, I can only give notice at the outset that I do reject those assumptions, and suggest that any viable theory of abstraction *must* reject them.

The theory I am going to defend was originated by Ayn Rand (1979), in a work that has so far received relatively little attention from philosophers. Her theory is quite different from any put forward by the classical British empiricists, but it shares with them a certain approach to the issues. 1) She assumes that the awareness of similarity, at least the mode of awareness involved in forming basic level concepts for perceivable objects, is perceptual. This is the way in which she escapes the circularity argument, and a large part of my exposition will be an attempt to describe this mode of awareness. 2) She is not a realist in either the Platonic or the Aristotelian sense. She holds that abstract properties and kinds do not exist as such—as abstract—apart from human conceptual processes. Concept-formation does not reveal abstract properties that exist independently in the objects themselves; it integrates and transforms the awareness of a network of similarities among the concrete, determinate properties of objects in such a way that we can grasp those properties in an abstract form, i.e., conceptually. 3) She holds that the process of abstraction can be segmented into two basic stages: grouping perceived objects into a class on the basis of similarities among them, and forming an abstract concept for the universal *type* to which those perceived objects belong. The first stage will occupy us for the next three sections; the final section will deal with the second stage.

Perception and Similarity

As I noted above, we can escape the circularity argument if the initial awareness of similarity is perceptual, not conceptual. Of course, this applies only to similarities among perceived objects, in respect of such perceptible properties as color, shape, weight, and the like. And it is only these properties, and the natural kinds that can be isolated on the basis of them, for which concepts are formed directly from perception. Even so, an initial difficulty besets the claim that these similarities are *perceived*.

One part of the problem is phenomenological. It seems evident, as

Cassirer (1953) points out, “that the ‘likeness’ between any contents [of the visual field] is not itself given as a further content; that similarity or dissimilarity does not appear as a special element side by side with colors and tones” [p. 16]. It is likely that Cassirer intended his argument to apply to similarity *qua* relation, since relations never appear as “further contents,” “side by side with” their relata. But that does not mean relations cannot be perceptually *given*. If we reject the idea of point sensations—and surely the Gestalt psychologists have shown that we should—then relations such as distance can be given in the same way as intrinsic features such as color. But Cassirer’s point does capture something peculiar to similarity. When two apples appear similar, there is nothing there in the visual field except the two apples, with the features each one has. What does it mean for them to appear similar? We can distinguish two cases: each of the apples can appear a certain way, and the appearances can in fact be similar, without the similarity itself being noticed; or the similarity may appear as well and be registered cognitively. This is certainly a difference in cognitive state, but it does not seem that there need be any difference in the perceptual field.

The other part of the problem is causal. A property of an object is perceptible when there is a causal link between that property and some feature of the perceptual field. In terms of Gibson’s (1966) model, we might say that in the normal case, some property *A* of the object’s appearance is causally specific to a property *S* of the stimulus array at the relevant sense organ, which in turn is specific to the property *O* of the object. In order to determine whether a given property *O* is perceptible, as distinct from some other property *O*’ (e.g., the mass of an object as opposed to its weight), we vary *O* while holding *O*’ and all other properties constant. If there are corresponding variations in the appropriate features *A* and *S*, which allow the subject to detect the variations in *O*, then *O* is perceptible. But there is no way to vary the similarity between two objects without also varying the properties in respect of which they are similar. It is also difficult to imagine a stimulus feature which is specific to the similarity between the objects, but distinct from the stimulus features specific to the properties in virtue of which they are similar. Thus the phenomenological fact noted above is precisely what the causal chain would lead us to expect: the similarity between two objects is given by the same sensory mechanism as the objects themselves, and the difference between noticing and failing to notice a similarity in the perceptual field does not (necessarily) involve any difference in the field itself.

That difference, rather, is a function of a higher-order process of *attention*. The physical nature of the external objects and of the sensory apparatus determine what is given perceptually, but attention allows us to

focus selectively within the field. The result of focusing, however, is still a form of perceptual awareness—of whatever one has focused on. Attention may be guided by concepts one has already formed, and it may result in a perceptual judgment which identifies conceptually the object attended to; in adult humans this is probably the normal case. But attending does not *consist in* conceptually identifying. Indeed, one could not form a perceptual judgment about an object unless attention had already isolated the object to be identified. Attention *per se* consists in heightening and sharpening the perceptual awareness of some object or aspect of the field. In reply to the circularity argument, then, we may still say that similarity is perceived, but in order to understand *how* it is perceived, we must recognize the crucial role of attention. A primary goal of a theory of abstraction is to describe the particular mode of attention that allows us to focus on similarity.

One strand in the empiricist tradition held that abstraction is a process of attending to what two or more objects have in common, and ignoring their differences. We will see that this is true in a very broad sense, but it is useless as a point of departure. Consider Locke's (1894) description of the process:

...let anyone effect, and then tell me, wherein does his idea of *man* differ from that of *Peter* and *Paul*, or his idea of *horse* from that of *Bucephalus*, but in the leaving out something that is peculiar to each individual, and retaining so much of those particular complex ideas of several particular existences as they are found to agree in. [III, 3, 9]

In schematic form, Locke's model is that we perceive particular objects (Peter, Paul) as complexes of elements: *abc*, *abd*, etc. The concept 'man' is then formed by isolating what Peter and Paul have in common (*a* and *b*), and disregarding what distinguishes them (*c* and *d*). Two obvious inadequacies of this account suggest two constraints that a theory of abstraction must accept.

If the result of Locke's process is to be the abstract concept 'man,' first of all, the properties *a* and *b* will have to be abstract ones—a rational faculty, an upright posture, the utterance of speech, etc. But we do not perceive abstract properties as such. When I see Peter, I do not see uprightness in the abstract; I see a determinate shape with a long vertical axis. Nor do I hear speech as such; I hear a determinate stream of temporally segmented tones and overtones. And at this determinate level, Peter differs from Paul in these respects no less than in the accidental ones such as skin

color (the *c*'s and *d*'s). Thus if I leave out of my concept 'man' the properties that are peculiar to Peter and Paul, I will have to leave out *all* the properties I perceive, and the concept will be empty. The problem here is that Locke simply ignores that real problem of abstraction, and it suggests our first constraint: a theory of abstraction should not simply assume, without explanation or justification, a capacity to grasp abstract features. For it is that capacity that we are trying to explain.

The explanation will make use of a capacity to notice the (determinate) similarity between Peter and Paul in respect, say, of posture or speech. The problem will be to explain how the awareness of that similarity gives rise to the conceptual awareness of the abstract respect in which they are similar. It would be helpful, then, if we could appropriate current psychological theories of similarity to serve in that explanation. Unfortunately, those theories are such that they cannot serve this function without violating our first constraint.

The current psychological theories can be distinguished by the relative emphasis they place on features or dimensions—the standard distinction being that features are properties which an object either has or lacks (*red, male, winged*), whereas dimensions (*size, weight*) admit of quantitative degrees (Smith and Medin, 1981). The best example of a featural theory is Tversky's (1977) model, according to which the similarity two objects are seen as having is a positive function of the features they share, and an inverse function of their distinctive features. This presupposes that each object is represented as a set of features, and the features will necessarily be abstract ones if the theory is to allow for a high degree of similarity between objects that differ minutely but perceptibly in each determinate property. For that reason, however, Tversky's theory presupposes the capacity to identify the abstract features possessed by the objects, and therefore will not serve our purpose.

Dimensional theories, by contrast, assume that each object is represented in terms of its determinate values among the dimensions that are relevant to the similarity judgment. Thus objects are represented not as sets of abstract features (*abc, abd,...*), but as sets of determinate ones ($a_1b_1c_1, a_2b_2d_1,...$). Similarity is then a summary function of the distance between the objects on each of the relevant dimensions (often represented as the distance in a metric space defined by the dimensions). As we will see, it is an essential part of Rand's theory that such quantitative relations do play a key role in the awareness of similarity. But we cannot simply incorporate any such "distance" model as currently formulated, because they all leave a crucial point unexplained: they assume the capacity to identify two particular properties, a_1 and a_2 , as determinate values of the same abstract dimension. In this more subtle way, they too violate our first constraint.

We shall have to look beyond psychology, then, for an adequate theory of similarity. (For a fuller statement of this argument, see Kelley and Krueger, 1984).

The second major problem with Locke's account (as well as with the psychological theories) is his assumption that the objects, as perceptually given, are already broken down into constituent properties—a problem that remains even if the properties are taken to be determinate ones (a, b, \dots). In Locke's case, the assumption is part of his sensory atomism, his view that the perception of an entity is a complex idea formed by composition from simple ideas of its various sensory qualities. There is ample reason to avoid such atomism (Kelley, 1986). Even though it is possible for an experimenter to determine from the outside what properties of an object the subject is perceiving, by varying the objects in appropriate ways, asking the subject to make matches, etc., nevertheless those properties may not be individuated *for the subject*. He is perceptually aware of different properties, but he may not be aware of them *as* different; individual properties may not stand out as figures against the background of other properties, in the way the object itself stands out against its background. One of the tasks of a theory of abstraction is precisely to explain how attention comes to be focused on particular qualities in such a way as to make them stand out.

Now the similarity between two objects will always be a similarity in some respect, such as color, which means: the similarity relation exists in virtue of the determinate properties of the objects, their particular shades of color. For the reason just given, we cannot assume the subject individuates the color property of each object separately before he is aware of the similarity (in respect of color) between them. For the same reason, then, we also cannot assume that the subject is capable from the outset of discriminating different dimensions of similarity. Smith and Kemler (1978) have found, for example, that young children make holistic similarity judgments about objects, judgments which are not simply the sum of judgments along component dimensions. When they are given objects that vary in both shape and color, they make judgments which treat those dimensions as a single, integral one. (This phenomenon is presumably part of the reason children typically acquire concepts for natural *kinds* of objects, which share clusters of correlated features, before they acquire concepts for the individual attributes in the clusters (Anglin, 1977).) In the same way, adults normally make similarity judgments about color as a unitary dimension, without attending selectively to hue, saturation, and brightness. As Smith and Kemler suggest, then, "The problem for a developmental mechanism or learning procedure is to make the dimension axes available and to endow them with a special experiential status" [p. 527]. What this suggests is that the awareness of similarity becomes increasingly differentiated, as the

subject discriminates different dimensions of similarity within the global relation; and that this awareness of particular dimensions of similarity is what makes it possible for the subject then to focus on the particular qualities in each object in virtue of which that relation holds. In any case, the second constraint is that we should not assume, without explanation or justification, the subject's capacity to individuate either particular qualities in objects or particular dimensions of similarity among objects.

(At some point, of course, we are going to have to posit basic capacities that cannot be further reduced to or analyzed in terms of more elementary capacities. Any theory of a cognitive process must commit itself somewhere to certain atomic subprocesses. My point here is merely that we must be able to justify the claim that we have reached that point in our theory; and that Locke's theory of abstraction did not in fact reach that point. I shall return to this issue in my conclusion.)

Abstraction: The First Stage

So much for the preliminaries: let us turn now to Rand's (1979) theory of abstraction. The first stage in the process is the grouping of objects on the basis of similarity: her view is that the essential achievement at this stage is the awareness of objects as 'units.'

When a child observes that two objects (which he will later learn to designate as "tables") resemble each other, but are different from four other objects ("chairs"), his mind is focusing on a particular attribute of the objects (their shape), then isolating them according to their differences, and integrating them as units into separate groups according to their similarities...

A unit is an existent regarded as a separate member of a group of two or more similar members.... Note that the concept "unit" involves an act of consciousness (a selective focus, a certain way of regarding things), but that it is *not* an arbitrary creation of consciousness: it is a method of identification or classification according to the attributes which a consciousness observes in reality [p. 7].

We will see in a moment why she uses the term "unit" to describe entities insofar as they have been grouped by similarity. The first point to notice, however, is that the similarity among objects which allows one to group them together is perceived in the context of their difference from dissimilar things. Tables stand out as similar against the background of their dif-

ference from chairs. This point alone will not take us very far, but it is important to be clear about exactly what it does say.

It is fairly common among similarity theorists to use a contrast object to define the relevant class of similar things, and the contrast object (or foil, as Quine (1977) calls it) is usually combined with a notion of comparative similarity: something is in the class if it resembles an exemplar *A* more than the foil *B* resembles *A*. The motivation for this is to avoid two problems with unqualified similarity: i) almost any object will be similar to *A* to some *degree*, however minimal, along any given dimension, and ii) any object will be similar to *A* in some *respect*, along some dimension. If we relied exclusively on unqualified similarity, we would have no way to explain why the child excludes chairs from the group he forms, since they are after all similar to tables to some degree even in shape, and are similar to tables in other respects as well. Now the comparative notion of similarity might solve the problem of degree, by requiring that any object to be included in the group be *sufficiently* similar to *A*. But it does nothing to solve the second problem.

The comparative notion seems to offer us just two choices. On the one hand, we could say the relation of comparative similarity holds, and is perceived as holding, among *properties* of the objects. Objects are grouped together as 'red,' for example, by grouping together all those objects whose color property is more similar to the color property of *A* than is the color property of the foil (something orange). But this would violate our second constraint, but assuming that the subject has the capacity from the outset to individuate properties of perceived objects. On the other hand, we could say that comparative similarity holds, and is perceived as holding, among *entities*; it would be expressed by a predicate taking objects as its arguments. This is the sort of relation that Carnap (1967) took as basic (though his wholes were not objects but *erlebs*, e.g., the entire visual field at a moment), and his motive was precisely to observe something like our second constraint. As Goodman (1977) pointed out, however, this second approach encounters the problem of "imperfect community": both *C* and *D*, as entities, may resemble *A* more than *B* does, without resembling *A* in the same respect. *A*, *C*, and *D* in that case will not form a class based on a consistent pattern of similarity. The pairwise similarity judgments—*A* and *C*, *A* and *D*, *C* and *D*, etc.—will be based on different dimensions of similarity.

The problem of imperfect community has a certain psychological reality. Goldstein (1940) studied a number of brain-damaged patients who lacked what he called "the abstract attitude." When asked to group together objects in the vicinity which "belong together with" a given object, they had great difficulty selecting more than a few objects along any single

dimension. Some of the patients could select objects consistently along such dimensions as shape, color, material, and so forth—if they were shown how by the experimenter, or if the relevant properties were made especially salient. But even then the groupings were rarely complete. More often, patients would switch randomly and inadvertently from one dimension of classification to another.

We observe...that the choice may be determined by a number of different attributes—at one time by brightness, at another by softness, or coldness, warmth, etc....Further, we observe that he does not seem able to hold to a certain procedure. He has chosen, for example, some bright skeins [of yarn]. Suddenly he begins selecting on the basis of another attribute [pp. 71-72].

These patients were able to make holistic similarity judgments. But they seemed to lack precisely that capacity for forming similarity groups which is required for the process of abstraction: the capacity to use the same dimension of similarity across all the pairwise comparisons of the objects at hand. Yet how can this capacity be understood without positing an unexplained capacity to individuate specific properties, or at least specific dimensions of similarity? The explanation which Rand's theory offers is based on the notion of comparative *difference*.

Rand is not claiming, in the passage quoted, that table *A* is seen as more similar to table *B* than the chair is; the tables are not seen as similar to the chair at all, but as *different* from it. Nor does she take for granted the capacity to focus selectively on the property of shape in each of the objects: the awareness of the tables' similarity in respect of shape is the result of the awareness of each table as different from the chairs in respect of shape. The similarity between *A* and *B* is brought out by their common difference from *C* (the chair), by the fact that they differ from *C* in the same way. The assumption behind this account is that the capacity to individuate and focus selectively on the respect in which *A* (or *B*) differs from *C* is more basic than and underlies the capacity to individuate the respect in which *A* and *B* are similar. More generally, individuating the ways in which things differ is the fundamental ability, which makes it possible to individuate the ways in which they are similar. She is also assuming a capacity to discriminate different degrees of difference (just as the comparative similarity theories assume the ability to discriminate degrees of similarity). What is the basis for these assumptions?

Consciousness, says Rand, is "an active process that consists of two essentials: differentiation and integration" [p. 5]. We will see as we go along

that at each stage of concept-formation we can distinguish integrating and differentiating aspects of the process, and these aspects are always complimentary. But given the perceptual basis of abstraction, differentiation has a certain primacy in the case at hand. The perception of an object is the discrimination of it from its background, the awareness of it as different from the things around it. Seen by themselves, two tables would be perceived as different: each would be part of the background from which the other is discriminated. Discrimination in turn requires the detection of differences between object and background in respect of at least some perceptible properties: a visual figure stands out from the ground because of a color difference at the edge, and, in three-dimensional perception, a difference in depth. Similarly for other properties in the other modalities. Thus the two tables differ perceptibly in at least some of their perceptible properties: color, shape, size, etc. The awareness of their similarity must overcome the awareness of their difference. Rand's theory is that the more radical difference between either one of the tables and a chair allows the subject to grasp the difference between the tables "from the other side," as it were: he sees that the two tables are not *as* different from each other as either one of them is from the chair. The awareness of the relation between the tables as a less-than-complete difference is the enabling condition for the awareness of them as similar. The primary notion in her theory is therefore not comparative similarity, but comparative difference.

This is still a very abstract description, however, of the structure of relations we attend to in isolating a group of similar objects. Rand goes a step further in describing that structure, specifically in describing how we detect different degrees of difference. When *A* and *B* are seen as differing less than either does from *C*, she says, *A* and *B* are seen as differing *quantitatively* in a respect in which both differ *qualitatively* from *C*. The difference in shape between the two tables is quantitative, as opposed to the qualitative difference in shape between either table and the chair. Correspondingly, the similarity between the two tables is essentially the fact that they are commensurable: "similarity, in this context, is the relationship between two or more existents which possess the same characteristic(s), but in different *measure* or *degree*" [pp. 15-16]. There is thus a deep connection between her concept of a unit as a member of a similarity group and the mathematical concept of a unit of measurement. This indeed is the fundamental innovation in her theory, and I want to examine it in some detail.

Similarity and Measurement

One of the major issues in recent psychological work on concepts has been the relative importance of features and dimensions as component

properties in concepts, especially in concepts for natural kinds of perceptible objects (Smith and Medin, 1981). As I indicated above, the same distinction is important in psychological theories of similarity. Its interest for us at the moment is that there is a correlation between the quantitative/qualitative and the feature/dimension distinctions, so that we may get a better understanding of the first by exploring the second. Features are said to be qualitative properties, such as *red* or *winged*, which an object either has or lacks; dimensions are quantitative properties, such as *size* or *weight*, which objects can have in different degrees. This apparent distinction implies or suggests three others.

1. If two objects are compared in respect of some feature, there is only one way in which they can differ: one has the feature and the other lacks it. We can find this same difference in the case of dimensions: one object can have a weight, while another (e.g., a rainbow) can be weightless. But objects can also differ in another way in respect of a dimension. They can differ in degree, and the difference is expressed numerically: an instance of the dimension serves as a unit, and the difference between that unit and any other instance can be specified numerically, with the same numerical system serving for all dimensions. Even when we do not know how to specify such quantitative relations as cardinal numbers, we can still rank the instances ordinally, and there is a common structure in all ordinal rankings. Thus in addition to the all-or-nothing differences that dimensions, like features, permit, dimensions also exhibit a structure of orderly quantitative differences.

2. Connected with that mathematical point is a logical one. It would be contradictory for an object both to have and to lack a given feature, and the same is true for dimensions. In the case of dimensions, however, it would also be contradictory for a thing to have two different values on the dimension. Contradictory predications are possible at two different levels: in the claim that something is both weighted and weightless, and in the claim that it is both three and six pounds. This is one reason why dimensions cannot be treated as sets of nested features, as some theorists propose (e.g., Smith and Medin, 1981; Tversky, 1977). In the usual conception, features are regarded as logically atomic; but two increment-features of weight would be logically related in a way that an increment-feature of weight is not related to an increment-feature of size.

3. Dimensions are typically more abstract and universal than features (Smith and Medin, 1981). When dimensions are used to specify conceptual contents in a certain domain, say *animals*, the content of each concept in the domain is specified partly as a range along the relevant dimensions: the contents of 'bird,' 'dog,' etc., can be specified partly in terms of the appropriate ranges of shape, size, etc. When features are emphasized,

however, each feature may be relevant to only one or a few of the concepts: *having wings* is relevant only to the content of ‘bird,’ *barking* to the content of ‘dog,’ etc. Of course, other animal species have the property of lacking wings (or failing to bark), but the difference between having and lacking such features is not *like* the difference between two values or ranges on a dimension, for the reasons given under (1) and (2) above.

These three points were offered as a way to spell out the standard distinction between features and dimensions, but as I hinted above, the distinction will not hold water. We can see why by looking more closely at each of the points.

1'. Features are abstract attributes, and therefore can be instantiated in different determinations. *Red* is typically treated as a feature, but it comes in many different shades. And the various shades of red differ along certain *dimensions*: hue, saturation, brightness. In psychological theory, relations among different values along these dimensions are specified in cardinal terms, with the “just-noticeable-difference” serving as the unit. But even if we deny that this is a genuine cardinal measure, and regard it merely as an ordinal ranking, the point remains that the differences among determinations of the abstract feature *red* consist in quantitative relations. Similarly, *flying* is normally considered a feature: either an animal flies or it does not. But the many determinately different ways of flying differ along such dimensions as rate of wing-beat, and ratio of beating to soaring (hummingbird vs. condor); speed of flight (crow vs. falcon); height (chickadee vs. hawk) and amount (chicken vs. Arctic tern) of flying. Again, an early childhood definition of food is “things you eat,” and edibility is usually considered a feature. But as with other functionally specified classes, foods vary along the dimension of how well they serve the function—how nutritious they are—as well as along such perceptible dimensions as sweetness. These examples show that a feature is after all a kind of dimension whose instantiations are related quantitatively.

2'. Features exhibit the same two levels of contradictory predication that dimensions do. A thing cannot be red and not-red, but neither can it possess two different shades of red (at the same time and in the same respect). A bird may be capable of two different types of flying, but a given act of flying cannot be characterized by two different speeds or altitudes. We also find the same logical pattern one level of abstraction higher. A thing cannot be red and blue at once, nor a given act of locomotion at once a case of flying and of walking. That is to say, features are not only dimensionally organized within their ranges of instantiation; features themselves are dimensionally organized along a higher-order dimension (color, locomotion). The same is true, of course, for natural kinds within a single genus: an animal cannot be both a man and a bear, because it cannot have

two different shapes, or lots of hair and little, or be capable and incapable of speech.

3'. As a consequence of (1') and (2'), finally, it is clear that the order of abstractness among features and dimensions is considerably more complex than suggested above. Two shades of red differ quantitatively from each other—a dimensional structure at the lowest level of abstraction. At the next level up, *red* is a feature that seems to differ qualitatively from *blue*, although *red* and *blue* can also be seen as values on the higher-order dimension *color*. Finally, the property of being colored or having a color seems to be a feature that contrasts qualitatively with being transparent or invisible.

Now let us look at this sequence from another angle. On Rand's theory the height of abstractness corresponds to—it is nothing other than—the width of differences that are disregarded. Thus if we start with two very close shades of red, the difference between them is naturally seen as quantitative. As we take objects that are farther and farther apart on the spectrum, there are points at which we encounter steeper gradients in difference of hue—points at which, as we move steadily across the spectrum, the rate of change of hue accelerates. The difference across such points is naturally seen as qualitative, as against the smaller differences among shades that lie between two such points. The latter shades are naturally grouped together, and seen as sharing a single feature. In the case of natural kinds, instances differ quantitatively along a number of relevant dimensions; at the borderline of a kind there are steeper gradients on all or many of the dimensions simultaneously.

Two points emerge from this discussion. First, the distinction between qualitative and quantitative differences is itself quantitative—a difference in the degree of difference. Second, the distinction is a relative one. When the difference between two objects is perceived, the properties of the objects do not determine whether the difference will be seen as qualitative or quantitative. That depends on the context in which that difference is perceived, in particular the context of other differences. In diverse contexts, the same difference may be seen as qualitative or quantitative. At one end of the scale, consider two minutely different shades of color. Most people would regard them as quantitative variants, if they noticed the difference at all, but it is possible to focus on the uniqueness of each shade in such a way as to make the difference between them appear qualitative. Good artists seem to have this ability. At the other end of the scale, the difference between colored and invisible objects normally seems qualitative, as I suggested above; but it would seem quantitative if opposed to the more radical difference between either sort of object and something non-physical, to which the colored/colorless distinction does not even apply.

(It is important to emphasize that I am speaking here only about the *epistemology* of our perception of difference. I am not making, nor would I defend, any metaphysical claim that all qualitative dimensions are ultimately reducible to some one or several basic dimensions.)

Berlin and Kay (1969) have provided an interesting example of this relativity in their cross-cultural study of color vocabularies. They found that basic color terms in the various languages were intertranslatable, in the sense that the speakers of different languages agreed on the focal instances of corresponding terms. They also found that basic color terms could be ordered as follows:

white	<	red	<	green	<	blue	<	brown	<	purple
black				yellow						pink
										orange
										grey

where $a < b$ means that any language which has b has a , but some languages possess a without b . Thus the color domain is partitioned differently, but not arbitrarily. Consider a language with just five primary color terms; these would have the same focal instances as our “white,” “black,” “red,” “green,” and “yellow.” In such a language, the word corresponding to “green” would be centered on the same shade that English speakers agree is pure, prototypically green. But the boundaries of the word would be quite different, since it would include a large segment of the shades we would call blue. On the one hand, then, such a language would treat as quantitative certain differences (between green and blue) that English treats as qualitative. That is not, apparently, because speakers of such a language cannot make the same *perceptual* discriminations we can, but rather because the difference between green and blue seems (merely) quantitative in comparison with the qualitative difference between green and yellow or red. On the other hand, the fact that the terms are ordered in the way they are suggests that there is something objective about degrees of difference, and thus of similarity, among the colors. It suggests that the difference between blue and green cannot be seen as qualitative until the larger difference between either shade and yellow has been identified as qualitative.

Another sort of relativity is worth mentioning here because it plays an important role in Rand’s theory. When two tables are seen as differing quantitatively in shape, as against a chair, the qualitative contrast with chairs is also in respect of shape. Qualitative differences no less than quantitative ones require *commensurability*: they are differences along a dimension. As she observes, “No concept could be formed...by attempting to distinguish long objects from green objects” [p. 16]. This means that the dimen-

sion of shape, on which tables are distinguished qualitatively from chairs, creates the potential for a *quantitative* comparison between them—a potential which is actualized at a higher level of abstraction, in forming a concept such as ‘furniture.’ We saw before that there is an alternation of qualitative and quantitative differences as we ascend the ladder of abstractness. This reflects the sequence of qualitative and quantitative differences as we increase the width of difference which conceptualization must overcome. The point I am making now is that the alternation is dynamic. Two shades of red differ quantitatively as against a blue, but red and blue come to be seen as differing (merely) quantitatively when they are jointly contrasted with a transparent piece of glass. The hierarchy of genus and species is a succession of qualitative differences reduced to quantitative ones (see Rand, 1979, Chap. 3).

According to Rand, then, the first mode of attention involved in abstraction, the kind of selective focus that allows one to group objects consistently on the basis of similarity, is essentially an ability to discriminate qualitative from quantitative differences along dimensions of perceived variation. This mode of attention involves two levels of differentiation: the awareness of differences among objects, differences that are given in perceptual discrimination; and the (higher-order) discrimination between qualitative and quantitative differences. Her theory that similarity is perceived as a less-than-complete difference explains why similarity is perceived as extending no further than it is—why the chairs are not seen as (even minimally) similar to the tables. And it explains how particular dimensions of similarity can be individuated, without presupposing the capacity to individuate aspects of an object’s identity in any non-relational way. These are all advantages of her theory, for the reasons given two sections ago.

When two objects are seen as similar, however, they are not merely seen as less-than-completely different. The awareness of similarity has an integrative element as well. We saw that commensurability is involved in perceiving both quantitative and qualitative differences. In the latter case, commensurability is perceived merely as difference: each table is perceived as different in shape from the chair. The quantitative relation between two tables, however, is perceived not merely as a kind of difference but also as a positive connection between their shapes.

The perception of table *A* by itself would involve the awareness of the particular shape characteristic it has, and the awareness of that shape would be self-contained. The same characteristic is grasped in a relational way, however, when *A* is seen as similar in shape to another table. Since the difference in shape between *A* and *B* is merely quantitative, *A*’s shape can be grasped as a quantitative variant of *B*’s. One could “get from” *B*’s shape to *A*’s by making nothing more than quantitative changes—length-

ening the top, shortening the legs, etc. This relational perception of a quality is analogous to the way we identify determinate properties in conceptual thought: we specify a unit of measurement, which fixes the dimension along which we are measuring, and then give the particular measurement of the object at hand. The unit is simply an instance of the dimension, an instance we have all agreed to use as a standard; and the measurement is the quantitative relation between the object and the unit, the very quantitative relation we are aware of in perceiving objects as similar. Thus Rand's view is that the relational method of identifying determinate properties is already present, in a primitive form, in the awareness of similarity. When similar objects are grouped together, they literally are units, in essence if not in explicit mathematical form. Each one could serve as a unit in terms of which the relevant property of each of the others could be specified as a measurement; and those determinate properties are seen in that relational way, *as* measurements.

Thus the quantitative relations among objects not only allow us to group them together on the basis of similarity; they alter the awareness we have of the objects so grouped. This integrative aspect of the awareness of similarity, we shall see, is the crucial link to the second stage of concept-formation.

Abstraction: The Second Stage

In philosophy as well as psychology, there is a certain give and take between cognitive structure and cognitive process. In psychology, both structures and processes are used in any experimental task, so that hypotheses about one cannot be tested independently of hypotheses about the other; and there is always the temptation to save an hypothesis of one type by making ad hoc adjustments in the other. A similar problem arises for us at the point we have come to.

The second stage in Rand's theory of concept-formation is the stage of abstraction proper, the cognitive process whose outcome is a concept. Now among the constraints to which any such theory is subject are those imposed by the nature of concepts themselves: a theory about a cognitive process must be sufficient to explain the cognitive structure resulting from the process. The more detailed our understanding of what a concept is, then, the more easily we can test any theory about concept-formation. Unfortunately, however, there is not very much we can say about the nature of concepts that is independent of our theory of concept-formation. Unlike a perception, for example, a concept is not a conscious cognitive state with a content that can be described phenomenologically. A theory about what a concept is, and a theory about how it is acquired, will thus have to be formulated together as a single package. Nevertheless, there are

a few assumptions about concepts that I think *can* be argued for independently of Rand's particular theory, and I want to begin by laying them out.

A concept is both universal and abstract. It is universal because it is open-ended: it subsumes a indefinite number of numerically distinct objects. A concept is abstract, on the other hand, because it subsumes an indefinite range of *qualitatively* distinct objects, objects which (as we can now say) differ in measurement. Thus the opposite of universality is particularity, but the opposite of abstractness is determinacy. Both of these features raise questions that a theory of concept-formation must answer.

Let us begin with universality. Even after I have grouped objects together on the basis of their similarity, I am still aware only of those particular objects, the two tables (say) that I happen to be looking at. My mode of awareness is still perceptual, and thus limited to the things that are present to my senses. But after I have formed the concept 'table,' and begin using it in conceptual thought (e.g., in grasping the fact that tables are solid objects), my thought is about *all* tables, everywhere. That is what it means to say that a concept is universal. Now there is no way to understand one's capacity for conceptual thought about a universal class of objects by simple extrapolation from the perception of a few similar objects; there is a difference in kind here. The reason is not merely that one could not possibly perceive all the tables there are. Even if one could, it would not be possible to hold all of them in mind, to attend simultaneously to all of them as units of a vast similarity class. There is a limit to the number of items one can attend to and deal with consciously at any given time. Hence the first mode of attention, which involves the awareness of each particular object as a distinct unit, could not possibly encompass more than a very small number of the instances of the concept.

If we are able to think universally about whole kinds of objects, then, it is because concepts serve a function that Rand calls "unit-economy." Psychologists who have noticed the limitation on the number of items one can attend to have often surmised that a key cognitive capacity is the ability to "chunk" information, so that more information can be handled without increasing the number of items we have to deal with (Miller, 1956; Anglin, 1977; Rosch, 1978). This is precisely the role of concepts. A concept is a new mental unit (in the ordinary sense of "unit," not the technical one defined above). A concept integrates an unlimited number of perceptual concretes: both those that actually served as perceptual units in forming the concept, and the unlimited range of others that could be included in the same similarity class. In Rand's words, "A concept substitutes one symbol (one word) for the enormity of the perceptual aggregate of the concretes it subsumes." In this way, "Conceptualization is a *method* of expanding man's consciousness by reducing the number of its content's

units—a systematic means to an unlimited integration of cognitive data” [p. 85].

The problem this poses for a theory of concept-formation is to explain *how*, by what process, the perceptual units in a similarity class are integrated to form the concept as a new unit. And the key philosophical dimension of the problem is set by the second feature of concepts: their abstractness.

Because it is abstract, a concept is a way of treating an entire class of objects as if they were identical. When I identify an object I encounter as a dog, I am identifying it just insofar as it is not different from any other dog, for I would have made a judgment with the same conceptual content had I encountered any other dog. My judgment does not express any perceptual knowledge I might have about the features of this dog that distinguish it from others. When I anticipate that it will bark, moreover, I do not anticipate that it will make any particular one of the many determinately different noises that fall within the abstract category ‘barking’; I do not anticipate, of this dog, anything different from what I would anticipate of any other. Thus a concept is a way of treating discriminable things as if they were identical. This has the advantage of filtering out a mass of information that is irrelevant to most cognitive tasks. More importantly, it is necessary if the concept is to function as a single unit, representing all its instances indifferently, and thereby expanding our cognitive capacity to deal with whole classes of objects we have never encountered.

The claim I have just made must be qualified somewhat. Recent psychological research, particularly by Eleanor Rosch (1978), has shown that we distinguish degrees of *typicality* among instances of concepts. Beagles are typical dogs, Pekinese are not; and there is a remarkable agreement among subjects in judgments of typicality. This phenomenon has been shown to have an important influence on various conceptual tasks such as recognizing new instances of a concept. And degrees of typicality correlate with patterns of similarity and difference among members of a kind. A concept does not turn its back entirely, therefore, on its origins in the awareness of such patterns. But typicality or representativeness must be distinguished from “instancehood” or membership in a kind. A beagle is a more typical dog than a Pekinese, but the predicate “dog” applies univocally to both; we are saying the same thing of each in calling it a dog, and the statement “*x* is a dog” has the same truth value for any instance *x*.

A theory of abstraction, then, must explain how we come to regard the units of a similarity group as identical. And this will mean explaining how we come to ignore the differences among them, the quantitative differences that were so important in perceiving them as similar. The first stage of the process involved a mode of attention that allows one to differ-

entiate a group of objects from their environment by focusing selectively but consistently on a particular dimension of similarity. Now we need a mode of attention that functions within the group so isolated, and allows one to differentiate the element of identity from the element of difference within the similarity relationships. But there is a problem here. The objects are *not* literally identical. Nor is there any apparent way to distinguish the elements. As Hume (1969) noted, “It is evident, that even different simple ideas may have a similarity or resemblance to each other; nor is it necessary that the point or circumstance of resemblance [the respect in which they are similar] should be distinct or separable from that in which they differ” [I, 1, 7]. Two patches of red are both similar and different in exactly the same respect, namely color. Unless we are willing to embrace metaphysical realism, the sameness and difference are not literally distinct elements within the color patches themselves. The awareness of them as the same and as different must therefore emerge from attention to different *relationships* among the patches. A theory of abstraction must describe this mode of attention by identifying these relationships.

Since I have been drawing parallels between Rand’s approach and that of the British empiricists, it may help in clarifying the problem to indicate why she does not consider nominalism an adequate solution. Insofar as she identifies the concept, as a new unit of cognition, with the word, I think she would agree with Berkeley and Hume in rejecting the idea, common to realists and conceptualists, that a concept is a cognitive state with a literally abstract content. The only objects of conscious awareness are particulars and words. But this shifts the burden of explanation—of explaining *how* concepts function abstractly—from contents to processes. The nominalists failed to discharge this burden. For one thing, nominalists have traditionally taken the awareness of similarity as a primitive in their theories, and have therefore run up against the problems I discussed in connection with the first stage of concept-formation.

But there are problems as well at the second stage, which nominalists also leave unsolved. A word functions abstractly: it is not a name for one particular, but refers indifferently to any member of a class. A word cannot be a collective name for the set of similar objects one has grouped by means of the first mode of attention. It is not a collective name at all: if “man,” for example, is a name for the set of all humans, then it would be false to say that x (some particular person) is *a* man. More importantly, the word refers indifferently to any member of an open-ended class, not merely to the instances that were at hand when one formed the concept. How then does the term come to be extended? Nominalists have sometimes spoken of an implicit rule: anything sufficiently similar to these objects is to be included in the class. But if concept-formation involves the literal use of a

rule, then the theory of abstraction *would* be circular, since a rule is already abstract. There must be an actual mental process by which the perceivable differences among the similar objects at hand are disregarded, so that any new instance can be assimilated to the new mental unit formed as a result of that process. Nominalists have not provided an account of this process. This is no more than the traditional objection to nominalism: that it does not explain *how* we can refer to triangles (say) just insofar as they are triangles without appealing to an abstract content *triangularity* (see Weinberg, 1965).

In order to understand Rand's solution to this problem, we should start from the fact of determinacy. The outcome of the first mode of awareness is a highly structured type of awareness—or rather, it is the awareness of objects as elements in highly structured patterns. But the content of the awareness is entirely determinate. I am aware of particular objects, and of their determinate properties; and I am aware of the determinate quantitative relations that exist among them. Yet the awareness of objects as similar does represent a step away from the simple determinacy of perception. When I perceive a red object by itself, I am aware of a color property which is in fact determinate. Insofar as my perceptual awareness is specific to the color, I could in a sense be said to be aware of its determinacy. But that could not be said in any full-bodied sense, because I am not aware of it *as* determinate. The color is there before me as a qualitative content. I can look at the color, or close my eyes and cease being aware of it; I can attend to the red, or to the blue patch from which I discriminate it. But neither of these contrasts is the contrast implicit in the claim that I see the color as determinate. Determinate—as opposed to what? I have no perceptual clue that there is any other way to grasp the color than by seeing it as I do now. When I see the similarity between this and another red object, however, I am in a more literal sense aware of its determinacy. For the determinacy of the object in respect of color is captured for me by the set of its quantitative relations with other red objects, the quantitative relations I am able to attend to by differentiating them from the qualitative differences between these objects and blue ones. To grasp the color property before me in the relational way described at the end of the last section, to see it *as* a quantitative variant of the properties possessed by similar objects, *is* to grasp that color property as determinate. It is to grasp its determinacy in a specific form.

Rand's theory about the second stage of concept-formation can now be presented quite briefly. The differentiating aspect or moment is to distinguish the specific measurements of each object, in relation to the others, from the fact of commensurability. Since the determinacy of each object is seen as a matter of its quantitative relations to others, we abstract from

determinacy by omitting or disregarding the specific measurements, and attending to each object merely *qua* unit. The integrating element of the process is the awareness of the dimension of similarity, the dimension along which the units are quantitatively related, as an attribute they share in different measure or degree. We are aware of the attribute as an axis on which each of the objects at hand, and an indefinite range of objects not present, can be given a place. In her own words:

If a child considers a match, a pencil and a stick, he observes that length is the attribute they have in common, but their specific lengths differ. The *difference is one of measurement*. In order to form the concept “length,” the child’s mind retains the attribute and omits its particular measurement. Or, more precisely, if the process were identified in words, it would consist of the following: “Length must exist in *some* quantity, but may exist in *any* quantity. I shall identify as ‘length’ that attribute of any existent possessing it which can be quantitatively related to a unit of length, without specifying the quantity.”

The child does not think in such words (he has, as yet, no knowledge of words), but *that* is the nature of the process which his mind performs wordlessly...

Bear firmly in mind that the term “measurements omitted” does not mean, in this context, that measurements are regarded as non-existent; it means that *measurements exist, but are not specified*. That measurements *must* exist is an essential part of the process. The principle is: the relevant measurements must exist in *some* quantity, but may exist in *any* quantity [pp. 12-14].

That principle expresses the insight which governs the second mode of attention in the process of concept-formation. It is not, of course, a rule which subjects consciously apply—otherwise the theory would be subject to the objection I raised against nominalism. The principle is only implicit in the way subjects attend to the units before them, and to the structure of relations they exhibit. It is implicit in the realization that the specific measurements can be disregarded, and that consequently any number of other objects, bearing any quantitative relation (within a certain range) to the units at hand, might be included in the group. In this way, the principle

explains both the universality and the abstractness of the concept which results from the process. It should be noted, however, that after the concept is formed as a new mental unit, the principle does function as something like a rule, in two respects: the possession of the concept involves a kind of mental set or readiness to omit the measurements of new instances as they are encountered; and to omit the measurements of new dimensions of similarity among instances, as they are noticed.

Attention is a conscious process. To what extent, then, are the structures and processes involved in measurement-omission conscious? This issue is too complex to be treated fully here, but Rand gives an indication of her answer:

When, in the process of concept-formation, man observes that shape is a commensurable characteristic of certain objects, he does not have to measure all the shapes involved *nor even to know how to measure them*; he merely has to observe the element of similarity.

Similarity is grasped *perceptually*; in observing it, man is not and does not have to be aware of the fact that it involves a matter of measurement. It is the task of philosophy and of science to identify that fact [pp. 16-17].

In the first mode of attention, the subject is aware of the units as similar to each other, and different from the contrast objects. The principle that the units differ (merely) quantitatively from each other, but qualitatively from the contrast objects, describes the structure of relations in reality which the subject experiences in the form of simple similarity and difference. The principle explains the experience of similarity in the way that, for example, the phenomenon of binocular disparity explains the visual experience of depth. The principle explains, among other things, why the similarity is perceived as extending no further than the class of units, since, from an external point of view, we could find *some* respects in which the units are after all similar to the contrast objects in *some* degree. In the second mode of attention, what the subject is aware of is that the units are the *same*, that they have something in common. The theory of measurement-omission explains this experience by explaining what it is about the objects themselves that allows the element of identity to be factored out from the element of difference. It explains how one can ignore their differences in a given respect without ignoring that respect altogether.

This latter formulation reminds us of Hume's puzzle. When we say that two objects are both similar and different, we normally mean that they are the same in one respect, but different in another. In the cases at hand,

however, the objects are similar and different in the same respect: two shades of red are both similar and different in color, two pencils in length, two tables in shape. Forming a concept requires some way to drive a wedge between the sameness and the difference. The realist impulse has always been to look for the wedge in the objects themselves: to locate some metaphysical difference, within the determinate characteristic, between two real constituents—one to explain the identity with other objects, and one to explain the difference. The first constituent is an abstract property or essence (*redness, animality*), the other a “determining note” or “individualizing condition” (see Coffrey, 1917, I, Chap. 9). Yet the determinate property itself, as it presents itself to perception, does not reveal any such partition. In Hume’s (1969) words again, “’tis evident at first sight, that the precise length of a line is not different nor distinguishable from the line itself; nor the precise degree of any quality from the quality” [I, 1, 7].

But if we consider the relations among objects, then Hume’s puzzle is soluble (though not without enriching his conception of the mind’s capacity for attending to relations). Metaphysically, the abstract attribute or kind that a concept identifies is really the set of determinate relationships among determinate characteristics that allow objects to be ordered quantitatively. The concept ‘length’ is abstract in that it names indifferently “that [determinate] attribute of any existent possessing it which can be quantitatively related to a unit of length” (Rand, 1979, p. 13). There is no abstract property *length as such* lurking behind these determinate lengths. Cognitively, a concept allows us to separate the way in which two objects are the same from the way in which they are different (even though they are the same and different in the same respect). It does so, not because attention has been drawn to any distinction within the thing itself, but because attention is capable of distinguishing two ways of regarding an object in the context of its quantitative relationships to other objects. The difference is perceived as a specific quantitative relation, the identity as the fact of mutual commensurability. In ignoring the difference between two determinate characteristics, we do not ignore the characteristics themselves (which would leave nothing), or their determinacy, or any “determining note.” We ignore their measurements, their quantitative relations to each other.

Rand’s theory does involve, however, a genuinely abstract mode of attention, one that goes beyond any purely perceptual awareness of determinate characteristics and relations, however complex. Otherwise, it would be faced with the following objection. In perceiving an object, we are aware of what is in fact a determinate characteristic. The perception of similarity between that object and another is said to be an awareness of a quantitative relation between them. Then we disregard or omit the quantitative rela-

tion. Why does that not leave us back where we started, with the awareness of the determinate characteristic again? Rand's view is that perceiving the quantitative relation and then omitting it is not movement from one object of awareness and then back again, but a kind of cognitive spiral: from the determinate characteristic to the quantitative relation to the characteristic as a value on a dimension defined by a set of relations. This requires that one differentiate the property of being quantitatively related as such, of having *some* measurements, from the particular measurements an object has. It is this abstract mode of attention that explains the abstractness of the concept being formed.

Conclusion

It may be helpful at this point to summarize, in schematic form, the major elements in Rand's theory. The numbered stages in the table below are component processes which together, in sequence, form the process of concept-formation. I think it would be wrong to demand *a priori* that a theory fit the pattern, evident in the table, of alternating steps of differentiation and integration. Yet the pattern does make sense: the awareness of difference in each mode isolates a content on which attention may be focused, setting the stage for the awareness of a more complex difference within that content in the next mode.

A theory of abstraction should be evaluated in light of considerations at two different levels. At the local level, the explanation that is offered for each component cognitive task should be consistent with what is known about the task, and should offer a genuine explanation of our ability to accomplish it. For each of the components in the table (except (1) and (2), which I have discussed elsewhere [Kelley, 1986]), I have already defended Rand's theory at this local level in the course of expounding it.

But there are also broader requirements we may reasonably impose on a theory of abstraction. One is that the components should fit together in a coherent way, making it intelligible how each stage leads to the next. I think there is no question that Rand's theory meets this requirement. A theory should also be general: it should capture regularities in the way different concepts are formed. Given the central role of similarity, this means meeting a challenge put forward by Fodor (1981): "It appears to be a brute fact that the ways in which things resemble each other don't much resemble each other. What is common to what cabbages have in common and what kings do?" [p. 275]. Rand's theory meets this challenge. In regard to the nature of similarity, I have tried to show by a range of examples that the pattern of qualitative and quantitative differences is a general one. As for abstracting from similarity, the rule implicit in the concept 'red' (an instance must possess some shade of red but may possess any), and the

rule implicit in the concept ‘cat’ (an instance must possess some ‘cat-like’ shape, size, etc., but may possess any)—these and other cases do exhibit a common pattern: an instance must possess some measurement, but may possess any.

Table 1

<i>Mode</i>	<i>Differentiating Element</i>	<i>Integrative Element</i>
Perception	1. Discrimination of object from background	2. Awareness of object as entity
Perceptual attention	3. Discrimination of individual ways in which entity differs from background	4. Awareness of (determinate) property of entity
Awareness of similarity	5. Discrimination of quantitative from qualitative difference, between object and others, in respect of a (determinate) property	6. Awareness of measurements (awareness of object as quantitative variant of other objects, in respect of that property)
Measurement-omission	7. Discrimination of specific quantitative relations from commensurability	8. Awareness of abstract property as dimension of commensurability

What about the constraints mentioned earlier, concerning basic processes? Before we answer that question, we need to consider the source and meaning of the constraints. One of the goals of any theory about a cognitive process is to analyze that process into its components. But analysis must end somewhere. Any such theory must posit certain basic components: basic processes that cannot be broken down further into cognitive subprocesses; and basic capacities (to perform the basic processes) that cannot be broken down further into constituent cognitive capacities. (The word “cognitive” here is a necessary qualification, since there will presumably be neurological and/or evolutionary explanations for the basic cognitive processes.) It is no objection to Rand’s theory, then, that she treats certain capacities as basic. The only question is whether the ones she has chosen are plausible candidates. More generally, we may ask: At what

point is it reasonable to conclude that one *has* identified a basic process or capacity?

This question has been hotly debated in the cognitive sciences. One common view, for example, is that cognitive processes must be decomposed until we reach basic operations that can be implemented on a Turing machine (see Dennett, 1978; Haugeland, 1978; Pylyshyn, 1980; and Searle, 1980). It would take another essay to explore this question fully. In this context, I can only say that I see no warrant for accepting that criterion for basic processes. I do not think there is *any* single, global criterion, any more than there is a single criterion for having reached a basic causal law in physics. Indeed, I do not think there is any criterion not already implicit in what I have said. A basic process must be general, for example. It would not be reasonable to posit a basic capacity for perceiving similarity in respect of redness *and* a basic capacity for perceiving similarity in respect of animality *and*...—a basic capacity should be general in regard to the various contents on which it is exercised. Beyond this one global constraint, the only further constraints are the sorts of local ones I mentioned above.

To spell this out in more detail, let us return to Rand's theory, beginning with the first stage of abstraction: the awareness of similarity. A theory concerning this stage should not assume, without explanation or justification, a basic capacity to individuate either particular (determinate) properties of a perceived object, or dimensions of similarity among perceived objects. Why not? Because there is good phenomenological and experimental evidence a) that perceived objects are not *given* as sets of distinct properties, but that the individuation of properties requires attention to the relations among objects; and b) that the ability to distinguish different dimensions of similarity is the result of a developmental process. Now Rand's theory is consistent with the constraint, because the first capacity appears on the table as #4, the second as #6, and both of these integrative capacities are explained by a preceding capacity for differentiation. Does it violate the spirit of the constraint, however, to treat these capacities for differentiation as basic? Not if we look to the reason for the constraint. The problem of imperfect community implies that we must have *some* discriminative capacity that is finer-grained than the capacity to differentiate the object as a whole from its background. The capacity that Rand's theory takes as basic—the capacity to discriminate different ways in which the object differs from its background—is the least addition one could make to the gross discriminative capacity involved in perception, and it is consistent with the evidence mentioned above to treat it as basic. The same may be said for #5, the capacity to discriminate degrees of difference.

In regard to the second stage of abstraction, the constraint was that one should not assume, without justification or explanation, a basic capac-

ity to grasp abstract properties. It would be an obvious violation of this constraint to say that subjects form an abstract concept directly, without intervening steps, from the perception of a single instance. Rand's theory certainly does not violate the constraint in that way. As we saw at the end of the preceding section, however, her account of measurement-omission does presuppose a basic capacity to differentiate the specific quantitative relations each unit bears to others, relations that are determinate and different for each unit, from the common abstract property of being so commensurable. Does this violate the constraint? Or can we explain and justify the claim that this capacity is basic? Two points here will have to suffice.

First, Rand has certainly narrowed the gap between determinate and abstract, from both directions. On the one hand, an object's determinacy in a given respect is manifested in a structure of quantitative relations it bears to other objects, and the awareness of these (determinate) relations focuses attention on the determinacy of the object, preparing the mind, as it were, to grasp the other side of the contrast—the abstract dimension of measurement. On the side of the abstract property or kind, abstractness in her view is nothing more than the dimension of measurement; and the possession of an abstract concept requires nothing more (though nothing less) than the awareness that units must have some place in an ordered structure of determinate relations, but may have any. Given this analysis of the gap, it is not unreasonable to think it is crossed by a basic capacity.

Secondly, I do not believe one can demand more than this under the second constraint without begging the question in favor of a radically reductive theory of concepts—i.e., without incorporating nominalism into the conditions for an adequate solution. If a concept is a distinct mental unit (and not merely a kind of collective mental name for an unintegrated set of particulars); if it is a way of regarding its instances as identical (not merely as similar), and thus of *disregarding* their specific measurements; then no matter how we decompose the process of concept-formation there must be a stage at which the awareness of determinate objects, qualities, and relations gives rise to an abstract awareness of them. For all the reasons given in the text, I think measurement-omission is that stage, and I do not see any way to decompose it further. Someday, perhaps, we will have a neurological explanation of the capacity for omitting measurements, and perhaps also an evolutionary explanation for our coming to have it. But I see no way to decompose it further into *cognitive* stages. As Aristotle (1941) said in a similar context, “the soul is so constituted as to be capable of this process” [II, 19, 100a13].

Acknowledgments

I am indebted to Allan Gotthelf, Roger Donway, Janet Krueger Andrews, Jennifer Church, and especially to Leonard Peikoff, for comments and criticism.

About the Author

David Kelley, Ph.D., is the founder and executive director of The Atlas Society. He is a nationally-known philosopher, teacher, and writer, who has taught philosophy and cognitive science at Vassar College and Brandeis University. His books include *Unrugged Individualism*, *The Art Of Reasoning*, *A Life of One's Own: Individual Rights and the Welfare State*, and *The Contested Legacy of Ayn Rand*.

References

- Anglin, J. *Word, object, and conceptual development*. New York: Norton, 1977.
- Aristotle. *Posterior analytics* (G. R. G. Mure, trans.). In R. McKeon (ed.), *The basic works of Aristotle*. New York: Random House, 1941.
- Berlin, B., & Kay, P. *Basic color terms*. Berkeley and Los Angeles: University of California Press, 1969.
- Carnap, R. *The logical structure of the world*. (R. A. George, trans.). Berkeley and Los Angeles: University of California Press, 1967. (First German edition 1928).
- Cassirer, E. *Substance and function* (W. C. Swabie and M. C. Swabie, trans.). Chicago: Open Court, 1923.
- Churchland, P. *Scientific realism and the plasticity of mind*. Cambridge: Cambridge University Press, 1979.
- Coffrey, P. *Epistemology*. London: Longmans, Green, 1917.
- Dennett, D. Artificial intelligence as philosophy and psychology. In D. Dennett (ed.), *Brainstorms*. Montgomery, Vt: Bradford Books, 1978.
- Fodor, J. *The language of thought*. New York: Crowell, 1975.
- Fodor, J. The present status of the innateness controversy. In J. Fodor (ed.), *Representations*. Cambridge, Mass.: MIT Press, 1981.
- Gibson, J.J. *The senses considered as perceptual systems*. Boston: Houghton-Mifflin, 1966.
- Goldstein, K. *Human nature in light of psychopathology*. Cambridge, Mass.: Harvard University Press, 1940.
- Goodman, N. *The structure of appearance*, 3rd ed. Dordrecht: D. Reidel, 1977. (Originally published 1951.)
- Haugeland, J. The nature and plausibility of cognitivism. *Behavioral and Brain Sciences*, 1978, 2, 215-60.
- Hume, D. *A treatise of human nature* (E. C. Mossner, ed.). Baltimore: Penguin, 1969. (Originally published 1739-40.)
- Kelley, D. *The evidence of the senses*. Baton Rouge: Louisiana State University Press, 1986.
- Kelley, D., & Krueger, J. The psychology of abstraction. *Journal for the Theory of Social Behavior*, 1984, 14, 43-68.
- Locke, J. *An essay concerning human understanding* (A. C. Fraser, ed.). New York: Dover, 1959. (Originally published 1690; original Fraser edition 1984.)

- Miller, G. The magical number seven, plus or minus two. *Psychological Review*, 1956, 63, 81-97.
- Pylyshyn, Z. Computation and cognition. *Behavioral and Brain Sciences*, 1980, 3, 111-69.
- Quine, W.V. Natural kinds. In S. P. Schwartz (ed.), *Naming, necessity, and natural kinds*. Ithaca, N.Y.: Cornell University Press, 1977.
- Rand, A. *Introduction to objectivist epistemology*. New York: New American Library, 1979.
- Rosch, E. Principles of categorization. In E. Rosch and B. Lloyd (eds.), *Cognition and categorization*. Hillsdale, N.J.: Erlbaum, 1978.
- Russell, B. *Problems of philosophy*. Oxford: Oxford University Press, 1959. (Originally published 1912.)
- Searle, J. Minds, brains, and programs. *Behavioral and Brain Sciences*, 1980, 3, 417-57.
- Sellars, W. Empiricism and the philosophy of mind. In W. Sellars (ed.), *Science, perception and reality*. London: Routledge and Kegan Paul, 1963.
- Smith, E. E., & Medin, D. *Categories and concepts*. Cambridge, Mass.: Harvard University Press, 1981.
- Smith, L. B., & Kemler, D. G. Levels of experienced dimensionality in children and adults. *Cognitive Psychology*, 1978, 10, 502-32.
- Tversky, A. Features of similarity. *Psychological Review*, 1977, 84, 327-52.
- Weinberg, J. The nominalization of Berkeley and Hume. In J. Weinberg (ed.), *Abstraction, relation, and induction*. Madison: University of Wisconsin Press, 1965.

The ATLAS SOCIETY

Objectivism, the philosophy originally developed by novelist-philosopher Ayn Rand, is a revolutionary body of ideas that challenges both the tired dogmas of the cultural Right and the empty relativism of the cultural Left. In their place it offers a new moral and social ideal. The essence of Objectivism is, in the words of Ayn Rand, “the concept of man as a heroic being, with his own happiness as the moral purpose of his life, with productive achievement as his noblest activity, and reason as his only absolute.”

Founded in early 1990, The ATLAS SOCIETY is now internationally recognized as the premier source for Objectivist research, analysis and development. Serving students, scholars and the public at large, the Center’s mission is to advance Objectivist ideas and ideals in our culture and society, and to create an open community of people who share those ideals. We are crusaders for rational individualism, and a home for rational individualists.

TAS offers a wide range of educational and social programs for the growing Objectivist community. These programs include research, training, public advocacy, courses, conferences, publications and social events. And because Objectivism is a philosophy of reason, not dogma, the Society conducts its work in a spirit of free inquiry and intellectual tolerance.

We at the Society believe that Objectivism offers the moral and philosophical basis for cultural regeneration. The ATLAS SOCIETY is laying the intellectual groundwork; finding principled solutions to the practical problems of modern living, and fostering a benevolent community rooted in respect for individual human life and liberty.

TAS’s membership program allows individuals to invest in the future of Objectivism. A not-for-profit 501(c)(3) educational organization, the Society relies on the support of its participants, through tuition and attendance fees, and on tax-deductible donations from individuals, businesses and foundations. For more information, write to The ATLAS SOCIETY, 1001 Connecticut Ave. NW, Suite 830, Washington, DC 20036., call 1-800-374-1776, or e-mail tas@atlasociety.org.

Explore our products and programs; participate in our events. Help us advance a revolutionary new moral ideal for the new century.

www.ATLASSOCIETY.org