METAPHILOSOPHY
Vol. 30, Nos. 1/2, January/April 1999
0026–1068

# WHY DO CONCEPTUAL ANALYSTS DISAGREE?[1]

HAROLD I. BROWN

**ABSTRACT**: The practice of a priori conceptual analysis requires that the concept being analyzed be available in the analyst's mind. The difficulties of analysis and the existence of disagreements among analysts are explained by distinguishing the implicit knowledge we have of these concepts from the explicit knowledge we seek. This view of disagreement assumes that those who disagree are typically attempting to analyze a single shared concept. In this paper, reasons are developed for replacing this guiding assumption with the alternative assumption that members of a conceptual community often have similar, but not identical, concepts. Adopting this alternative requires that we rethink the nature and aims of conceptual analysis.

Key words: guiding assumption, conceptual analysis, conceptual community.

## Introduction

According to a widely held view, when philosophers analyze a concept they are seeking an explicit account of the concept's content – a content that they already know in some implicit manner. This implicit knowledge provides the intuitions that guide us in formulating proposed analyses, and allows us to recognize counterinstances to these proposals. Our inability simply to state the correct analysis is explained by this distinction between the implicit knowledge we already have and the explicit knowledge we seek. The same distinction is invoked to explain the wide-ranging disagreement on the analysis of familiar concepts that pervades the philosophical literature. It is a basic presupposition of this literature that philosophers typically disagree about the proper analysis of a single concept. Moreover, it is generally held that the concepts philosophers seek to analyze are shared by members of a wider community, although views differ on the extent of this community. While some hold that it includes all cognitive agents, others restrict its scope in specific ways. For example, some restrict common concepts to a culture or subculture, others to a

[1] I want to thank Herman Stark and two anonymous referees for helpful comments on an earlier version of this paper.

historical period of a culture, and others to a discipline or a historical period of a discipline. We might go further and suggest that concepts are idiosyncratic and that conceptual analysis is a form of autobiography, but this view is not commonly found among those who pursue conceptual analysis. Indeed, such a view would eliminate the disagreements that, to a large extent, drive the literature. In other words, the following proposition (I will call it *SC* for "same concepts") is central to the project of conceptual analysis as usually understood:

*SC*: All members of a conceptual community share a set of concepts.

My aim in this paper is to challenge *SC* – although not in a way that makes concepts idiosyncratic. This challenge will lead to an alternative view of the task of conceptual analysis and of what it is that binds a conceptual community together. Before embarking on this challenge, I want to clarify several points about *SC*.

1. I use the term "conceptual community," rather than the more common "linguistic community," in order to avoid any assumptions about the relation between language and concepts, such as whether nonlinguistic animals have concepts. However, I will be concerned only with humans in this paper and will follow the standard practice of using linguistic evidence as evidence about underlying concepts.

2. *SC* is not just a stipulative definition of a conceptual community. In accordance with customary philosophical practice, I will assume that members of a conceptual community can be recognized on the basis of their linguistic (and perhaps other kinds of) practice. For example, it is commonly assumed that when two people use the same word, this is a prima facie indication that they share a concept. It is also recognized that this is not always the case and that exceptions run from obvious ambiguities to more subtle variations. However, except in the case of obvious ambiguities, the burden of proof lies with those who would challenge the claim that shared linguistic practice indicates shared concepts. Pending further discussion, I will adopt the usual practice, although I note that one of my aims in this paper is to argue that subtle variations are more common than is usually recognized.

3. I will treat concepts as mental entities. The significance of this point can be brought out by considering Putnam's distinction between concepts and meanings (1975b, 217–18, 226, 245, 248). Concepts, Putnam maintains, are in the head, while the meanings of linguistic expressions are not in the head – so meanings are not concepts. Recognizing the existence of concepts is vital for understanding key features of the development of knowledge. To see why, let us grant for the sake of discussion that all members of a natural kind share an essence and that the discovery of these essences is one task of scientific research. At a given stage in the development of science we may not yet know the essence of a particular natural

kind, although researchers who are attempting to discover that essence will have formulated hypotheses as to its nature. In a familiar sense of the term "concept," researchers have a concept of this natural kind in mind, and this concept serves as the basis for thinking and communicating about that essence. Indeed, once we focus on the process by which we discover essences of natural kinds, it becomes a commonplace that our attempts to think and communicate require concepts understood as mental entities. Consider, for example, a stage in the development of chemistry at which modern notions of chemical structure have not yet been developed, so that researchers do not think of the essence of water in terms of its chemical formula. People are able to pick out examples of water, and some of these people are trying to understand the nature of water and to explain its role in physical and biological processes. In doing so, they will formulate a concept of water – perhaps as the cold, wet element – and an adult who has learned the concept while growing up in the society will be able to arrive at this account of water just by reflection. In other words, early chemists may have had an incorrect concept of water, but they had a concept nonetheless. The project of discovering the essence of a natural kind amounts to seeking a concept that accurately describes that essence. At any stage in the course of this research, the meaning of the label we attach to a natural kind is determined by that concept – not by the unknown (and possibly, at the time, inconceivable) essence.[2]

We get further support for this point by considering concepts such as phlogiston, caloric, and the Aristotelian notions of natural and violent motion. These concepts do not pick out natural kinds in this world, but there were people who thought about aspects of the world in terms of these concepts and who could have provided analyses of them. Twentieth-century folk with an interest in the history of science can still learn these concepts and provide such analyses. Moreover, putative instances of these concepts can be picked out in a nonarbitrary, intersubjective manner. Thus Aristotelians agreed that a falling stone is an instance of natural motion, and we can understand what they meant by this claim and why they believed it. The key difference between us and our Aristotelian ancestors is that they believed that the concept in question is instantiated in this world while we reject this claim. Note especially that we must acquire a concept before we can have an informed belief that it is not instantiated.

4. Some philosophers consider concepts to be abstract entities, but this view need not conflict with the thesis I am defending. For even if concepts are abstract entities, it is the images of these entities in individual psyches that generate our intuitions. *SC* can then be read as stating that the same

[2] Some philosophers distinguish concepts from "conceptions," but the central role that the entities I am concerned with play in thinking about a topic guarantees that the differences I am discussing are differences in concepts. Some will ask whether *all* our beliefs about a natural kind contribute to meaning of the associated term. Fortunately, we need not attempt to settle this question here.

abstract entities are implemented in the minds of the members of a conceptual community. The situation is analogous to that of a computer program, which is an abstract entity, but is implemented in a variety of different machines.

A similar point holds with respect to the view that concepts are ultimately social entities. Detailed discussion of this view is beyond the scope of the present paper, but whatever its merits, common analytic practice requires that each individual analyst have access to some entity available in her own mind, even if this is an image of a social entity.

Graham and Horgan propose that we recast our understanding of analysis as a "broadly *empirical* enterprise" (221) in which "our own introspectively accessible linguistic intuitions" (222) provide defeasible evidence about our concepts. Reflections by other members of the community provide one source of defeaters. It should be clear that this proposal also requires entities located in the individual psyche that are responsible for each individual analyst's intuitions.

5. Two other ongoing disputes are worth noting briefly. One concerns holistic versus atomistic accounts of meaning – and thus of conceptual individuation (cf. Fodor and LePore Chap. 1). The argument in the present paper will be neutral with respect to this dispute. The second issue concerns whether each concept is individuated by a set of necessary and sufficient conditions, or whether concepts exhibit "open texture." The latter view was developed in Wittgenstein's later work and has received considerable support from recent psychological literature. (See Lakoff Part I for a review of the literature.) It has also been argued that if this view is correct, it will require substantial changes in the practice of conceptual analysis (Bishop, Ramsey). I will develop the argument of this paper in terms of the necessary-and-sufficient-conditions view of concepts because this view is held by the overwhelming majority of conceptual analysts. However, the results of the discussion will also apply to any version of an open-texture view that makes sense of the notion of two individuals (or one individual at different times) associating the same concept with an expression. Open-texture theories will require different criteria for individuating concepts than necessary-and-sufficient-conditions theories, but those who hold that concepts are open textured do not automatically relinquish any notion of concept individuation.

6. We have intuitions about both the intension and extension of our concepts, and the practice of analysis requires that we bring these into equilibrium. As we pursue this equilibrium, we reconsider both proposed analyses and intuitions about whether actual or imagined items fall in the extension of a concept. Philosophers analyzing a concept may disagree on either or both; we will encounter examples of these situations in the next section. Conceptual analysts also disagree on whether intensions or extensions are fundamental in determining the content of a concept, or whether these are equally important, but I need not take a position on the ultimate

structure of concepts here. Three theses are central for present purposes: (1) philosophers seek analyses of concepts that are held to exist in their minds; (2) these concepts generate intuitions that provide the basis for analysis; and (3) analysts who disagree are typically attempting to analyze the same concept.

**Guiding Assumptions**

Since I will be arguing that *SC* should be replaced, it is important that we be clear on just what kind of proposition *SC* is. The obvious starting place is to ask whether *SC* is analytic or synthetic, but it will be more useful to begin by recalling a specific type of challenge to the analytic/synthetic distinction. One version of this challenge occurs in Putnam 1975a, where he argues that the distinction is not exhaustive. To see the point of Putnam's challenge recall that we are concerned only with universal propositions. Putnam notes that the characteristic feature of a traditional synthetic universal proposition is that it can be refuted by isolated experiments (e.g., 37, 46, 48); "All swans are white" is a classic example. Analytic propositions are immune from any such observational challenge. But, Putnam argues, there are also propositions that are not analytic, since they are subject to empirical challenge, but are not synthetic in the traditional sense because they are protected from refutation even in the face of evidence that, logically speaking, could count as a counterinstance. "These statements, then, have a kind of preferred status. They can be overthrown, but not by an isolated experiment" (46).[3] For example, the claim that energy is conserved in all physical interactions has faced numerous empirical challenges in the history of science. In most of these cases scientists protected the principle from refutation by postulating new forms of energy that had to be included in the accounting, although with the advent of special relativity and the recognition that mass and energy are interconvertible, this principle was incorporated into a more general conservation principle. None of the moves that occurred in order to protect energy conservation were arbitrary. Contrary evidence was not ignored; rather, the reluctance to reject this proposition in the face of apparent counterevidence led to new research on the newly proposed forms of energy.

Putnam argues that propositions in this third class are a heterogeneous mixture, but they share at least one common feature: They all have a degree of systematic import in our overall conceptual system so that rejecting one of them has consequences for our understanding of the world that go beyond the replacement of a simple empirical generalization. Some of these propositions have greater systematic import than others, but

[3] The passage continues: "They can be overthrown only if someone incorporates principles incompatible with those statements in a successful conceptual system." I will pick up this theme below.

in all cases, our reluctance to reject them has the effect of guiding research in a way that analytic propositions cannot do.

Putnam's third type of proposition is a variation on a theme introduced into philosophy by Kant. For each of Kant's synthetic a priori propositions we can provide consistent descriptions of empirical situations that, from the point of view of formal logic, are counterinstances to that proposition. We may, for example, measure the angles of a triangle and find that they do not sum to 180 degrees. If the proposition "The interior angles of every triangle sum to 180 degrees" is analytic, we must conclude either that there was an error in the measurements or that the three points do not instantiate a Euclidean triangle. Since Kant took it as given that Euclidean geometry is instantiated in the physical world, the only appropriate conclusion is that the measurements were not done with sufficient care. A similar account holds for all the synthetic a priori propositions on Kant's list.

Kant maintained that refusing to reject a synthetic proposition in the face of apparent counterevidence can be rational only if that proposition is known a priori, and he attempted to provide a justification for this status. We, however, can accept Kant's point that some propositions have this special status in research without following him down the path to transcendental logic. Propositions of Putnam's third type play the same methodological role as Kant's synthetic a priori propositions without being known a priori, and they can be rejected under appropriate circumstances. The central role of such propositions has become a familiar theme in contemporary philosophy of science. Sellars's (1963, Chap. 10) propositions that are nonanalytic but true *ex vi terminorum* are of this sort, as are Toulmin's "Ideals of Natural Order." It is a major function of Kuhn's "paradigms," Lakatos's "hard core," and Laudan's "research traditions" to provide such propositions.[4] In Laudan et al. the term *guiding assumption* was introduced to describe propositions that play this role. This term has the advantage of capturing a major function of these propositions while being neutral with respect to the various extant accounts of their nature; I will adopt this term to describe propositions in Putnam's third class. Antony has argued that the need for guiding assumptions is characteristic not just of science, but of cognition in general. This thesis will provide the starting point for the main argument of the present paper: I will argue first that *SC* plays the role of a guiding assumption for conceptual analysts, and then that *SC* should be replaced.

The existence of major disagreements among conceptual analysts makes it clear that *SC* is neither analytic nor a traditional synthetic proposition. *SC* is not analytic because disagreements could consistently be viewed as empirical evidence that *SC* is false. On the other hand, if *SC* were a traditional synthetic proposition, the massive disagreements that

[4] For a recent discussion see Brown 1996.

characterize the literature of conceptual analysis would require giving serious consideration to the hypothesis that *SC* is false. Instead, the usual response is to maintain *SC* and conclude that the fault lies with individual conceptual analysts so that, at best, all but one of those who disagree must have gotten the analysis wrong.[5] In other words, for each concept, *SC* directs researchers to seek a single analysis that will withstand all relevant counterinstances.

It will be useful at this point to recall just how wide a range of disagreement we find in the case of some of the central concepts that philosophers regularly examine. Consider, first, the concept of a causal relation. Philosophers disagree (among other issues) on whether it is conceptually possible for a cause to follow or be simultaneous with its effect (e.g., Brand; Hume 75–76; Kant 227–28; Kline; Mackie Chap. 7; Taylor 35–39); whether the relation is asymmetric or nonsymmetric (e.g., Mackie Chap. 7; Salmon 147–57, 158–68; Suppes Chap. 2; Tooley 178–80),[6] and even whether it is transitive (Lee); whether the concept of a cause includes (in some important sense) a necessary condition, a sufficient condition, both, or neither (e.g., Mackie Chap. 2; Salmon 185–90; Tooley 252–54); what the proper causal relata are (see Menzies for a survey); and what relations (if any) hold between causality and determinism (e.g., Mackie 190–91; Miller 39–43, 69–72; Salmon 122–23, 189–90; Suppes 6–8). Reflection on such disagreements has led Kim (112) to question whether there is "a unitary concept of causation that can be captured in an enlightening philosophical analysis."

Similar examples can be provided by examining disagreements over the analysis of such concepts as perception, truth, good, art, and just about any concept to which philosophers have turned their minds. The case of knowledge is particularly interesting because before 1963 there was general agreement that knowledge is justified true belief. That agreement dissolved in the face of Gettier's counterexamples, and at present the analysis of knowledge is as much in dispute as the other concepts just noted. (I will return to the analysis of knowledge in the following section.) Each of these cases can be considered evidence against the thesis that the disputants all associate the same concept with the familiar term. Widespread reluctance to consider this option underlines *SC*'s status as a guiding assumption. In other words, *SC* has a status analogous to that

[5] Compare Kuhn's remark that, in periods of normal science, failure to solve a problem within the constraints of the paradigm "is usually just a research failure, one which reflects not on nature but on the scientist" (35).

[6] A comment is in order on the relation between temporal and causal asymmetry. Anyone who requires that a cause precede its effect in time must also hold that the causal relation is asymmetric. But the relation between temporal and causal relations is a contested issue. Mackie and Tooley both insist (for different reasons) that causal relations are independent of temporal relations, but Mackie holds that causality is nonsymmetric while Tooley holds that it is asymmetric.

which is held, or once was held, by such claims as that all celestial motions are circular, that energy is conserved in all physical processes, that it is an objective fact whether two events that occur at a distance from each other are simultaneous, and many other central scientific claims. These are all examples of empirical claims that have been considered immune to empirical challenge and that played a guiding role in research for substantial periods of time. However, they are also all claims that are subject to reconsideration under appropriate circumstances. No one currently has an adequate general account of what these circumstances are, but it is still possible to make a case for abandoning a specific guiding assumption in a specific context. Moreover, one requirement for abandoning a guiding assumption is reasonably clear: An established guiding assumption is not rejected until an alternative guiding assumption that can take over its methodological role has been proposed, and reasons provided for thinking that the alternative will generate a more fruitful research program. I will now propose and defend an alternative guiding assumption for conceptual analysis.

## Conceptual Variation

I will call my alternative proposal *CV* for "conceptual variation":

> *CV*: Members of a conceptual community often associate different, although similar, concepts with a given term.

Obviously, *CV* can account for the existence of disagreement among analysts. *CV* allows for the possibility that a larger number of analysts are providing correct analyses than does *SC*, although these are not analyses of exactly the same concepts. However, this point is unimpressive as support for *CV* since *CV* was explicitly tailored to provide this result. Serious support for *CV* requires independent reasons for accepting it, responses to challenges, and reasons for thinking that its adoption will provide a basis for fruitful research. Note especially that we should expect research in a discipline to undergo some changes when we replace one of its guiding assumptions.

### *Refuting Challenges to CV*

I will begin my defense of *CV* by considering two challenges: We must explain why people who are using different concepts are still able, by and large, to communicate; and we must account for the common illusion among analysts that they are disagreeing about a single concept. Both challenges can be handled by the thesis that members of a community deploy similar concepts; let us explore the nature of this similarity.

The distinction between a concept's intension and extension provides two dimensions on which concepts may be the same or differ. Moreover,

we may find substantial intuitive agreement on one of these dimensions along with disagreement on the other; or we may find limited clashes of intuitions on both dimensions. In the case of the causal relation, much dispute occurs on the level of intension while there is substantial agreement on the extension. For example, Mackie holds that it is conceptually possible for an effect to precede its cause (Chap. 7), but this does not generate any intuitive clashes on specific cases with those who consider backwards causality absurd. Since Mackie holds that time-reversed causality never occurs in the actual world, such clashes will not appear with respect to any actual cases. Mackie does give an extended discussion of a small number of hypothetical examples that would provide evidence of backwards causality (173–92), but he works hard to try to make these examples plausible. They are not offered as examples that anyone would intuitively recognize as instances of backwards causality in the way Gettier's examples are offered as intuitively clear cases of justified true beliefs that are not knowledge.

A similar situation arises with respect to disagreements over the causal relation's symmetry properties. Differing views on the causal relata provide one source of such disagreements. For example, Suppes maintains that causal relations hold between event tokens so that if *A* caused *B*, it follows trivially that *B* cannot cause *A*. Others follow Hume in holding that causal relations occur between event types. On this view, that *A* causes *B* need not imply anything about whether *B* can cause *A*. Yet this dispute need not involve conflicting intuitions on whether any specific case is a clear instance of a causal relation, even though we are being offered significantly different analyses of this relation. In a similar way, philosophers can disagree on whether truth should be analyzed as correspondence, coherence, or disquotation (among other options) without any substantive dispute on which propositions are true – although some of the more extreme pragmatic and relativist accounts of truth will yield disagreements on this level. Note also that disagreements on a concept's intension may have only limited scope. For example, philosophers who disagree over which, if any, temporal relations are implied by a causal relation can agree on whether the relation is symmetric, transitive, and irreflexive. And philosophers who disagree on symmetry may still agree on the latter two properties just mentioned.

In other cases, intuitions clash with regard to both the intension and extension of a concept. In the case of causality, von Wright not only maintains that it is conceptually possible for an effect to precede its cause, but thinks that he can identify actual instances in which this occurs (76–77). Von Wright maintains that a case of this sort occurs when I raise my arm: There is a set of neural events that constitute necessary and sufficient conditions for this action, and that occur before the arm is raised, but these neural events are caused by the act of raising my arm. From the perspective of conceptual analysis the important point is that von Wright consid-

ers this to be a clear instance of backwards causality; the fact that he offers an actual example is irrelevant in the present context, since an imaginary example would have done just as well. Philosophers who maintain that it is conceptually impossible for an effect to precede its cause will not share von Wright's intuitions about this example. Yet, once again, the disputed cases are limited to a small subset of the overall set of causal interactions; for the most part von Wright's intuitions about the temporal direction of a specific causal relation will agree with the intuitions of those who consider backwards causality to be conceptually impossible. Much greater disagreement on both intension and extension occurs among those who offer analyses of the concept of art, and in this case communication does, sometimes, break down. A similar situation arises in contemporary disputes about the concept of a human life, although in this case disagreement on the extension of the concept centers on whether one especially important case constitutes a clear instance.

It will be illuminating to compare these cases to a situation that arose in the development of astronomy. Pre-Copernicans and post-Copernicans disagreed on both the intension and extension of the concept of a planet, although the scope of the disagreement was considerably greater with respect to intension. Disagreement on the extension was limited to whether a specific identifiable object, Earth, should be included in the class containing Mercury, Venus, Mars, Jupiter, and Saturn, or considered the unique member of a distinct class. Imagine two conceptual analysts, one living before the Copernican revolution, and one after this revolution was completed. Their intuitions about the intension of the concept of a planet would be drastically different, while their intuitions about its actual extension would have only minor differences. Now views on whether a major conceptual change occurred in this case vary with theorists' views on the relative importance of intension and extension in determining the content of a concept. For example, philosophers such as Feyerabend (1962, 1975) and Kuhn, who argue for drastic conceptual change, also hold that intension (analyzed in terms of inference patterns) is the only important dimension for comparing concepts. Others, such as Kitcher and Scheffler, who argue that the change was less drastic, consider extension to be the only important dimension. It would seem that these disputants are employing different concepts of a concept, but doing so in the context of substantial agreement about which cases in the history of science are especially relevant to the debate.

Cases in which there is agreement over a concept's intension but disagreement over extension also occur. Even if intension determines extension, we should expect such differences as long as a concept's intension does not provide an algorithm for picking out instances. For example, philosophers of science of an instrumentalist bent typically maintain that there is no conflict between classical mechanics and relativity since these are just alternative computational techniques that apply in different

domains. These philosophers are employing the concepts of a *classical entity* and of a *relativistic entity.* The intensions of these concepts are clear, but their domains of application shade smoothly into each other so that situations arise in which it is not clear whether a specific item should be viewed as classical or as relativistic. Van Fraassen (15–17) makes a similar point in defending the distinction between *observables* and *nonobservables* from Maxwell's critique. Maxwell had attacked the distinction on the basis of a continuity argument and van Fraassen replies that the two concepts (i.e., the two intensions) can be clear even though there is a gray area in which it is difficult to decide which concept applies. Moreover, outside of this gray area there is general agreement in our intuitions on which items are observable and which are not.

My proposal, then, is that philosophers who are debating the analysis of a concept often have in mind concepts that are very similar, but not identical. Sometimes intuitions clash only when discussion turns to extreme or unusual situations. If these intuitions are generated by the underlying concepts, then we have a strong reason for suspecting that the concepts in question are not identical. Still, the overlaps in proposed analyses and paradigm instances provide the basis for discussion. In a similar way, the general success of communication in a community can be explained by the large degree of overlap that we find; failures of communication will occur when these overlaps fail, and the existence of such failures supports *CV*.

There is an analogy between the kind of overlap I am discussing and Wittgenstein's notion of a family resemblance. In Wittgenstein's account, the family resemblance applies to instances of a concept and results in agreement over instances even though there is no set of necessary and sufficient conditions that all these instances share. Instead, the class of instances is bound together by similarities among specific cases. As I noted in the first section, I will not defend any position on the Wittgensteinian view of concepts in this paper. But whatever the structure of an individual's concepts, *CV* maintains that members of a conceptual community often have concepts that are related to each other only by a family resemblance – which may occur among both intensions and extensions. These family resemblances, not a set of identical concepts, bind the community together.

This idea of similar concepts in a community should be carried one step further: Individuals sometimes associate more than one concept with a single term. For example, those who have studied relativity physics have acquired two "time" concepts that they use in different contexts, and they can easily shift from one of these to the other as the situation requires. In this case we are not dealing with a simple ambiguity such as occurs with a word like "wound." The everyday and relativistic time concepts have many strands in common; these common strands provide the basis for teaching the relativistic time concept to someone equipped initially with only the everyday concept.

There is one difference between *CV* and *SC* that may appear to provide another objection to *CV. SC* has the advantage of being a truly universal proposition so that it indicates how to proceed in all cases of disagreement. *CV* only indicates an option that we should attend to, but leaves open the possibility that two philosophers who disagree on a specific analysis are discussing the same concept. Moreover, in this respect, *SC* may seem more like classic examples of guiding assumptions from the history of science, such as those which direct researchers always to reduce celestial motions to circles, or always to seek a force when an object is not moving at constant velocity. But not all scientific guiding assumptions operate in this straightforward way. For example, while Ptolemaics and early Copernicans agreed that all celestial motions are circular, and Kepler maintained that the planetary orbits are all ellipses, with Newtonian mechanics we give up any guiding assumption that states the shape of planetary orbits. Instead, we have a set of principles for calculating these orbits, and these new techniques generally do not yield a simple geometric shape. Moreover, these new guiding assumptions do not have the ease of application of some of their predecessors. Newton's second law, for example, is actually a law schema that must be filled in in specific cases. Much of the creative work involved in applying Newtonian physics consists of figuring out how to apply this law. A similar point holds for what is often referred to as *the* Schrödinger equation and for many other guiding assumptions, such as the principle that the number of atoms of an element in a molecule is a small whole number. The important point is that guiding assumptions reflect what we believe to be true about the domain in question – concepts, in our case. There is no guarantee that our best-established beliefs about a domain will be expressible in simple universal propositions.

I submit that the uncertainty involved in *CV* is an accurate reflection of the actual situation that conceptual analysts face. Sometimes analysts who disagree are analyzing different concepts; but sometimes they are examining the same concept and some of them have not yet achieved an accurate analysis. Moreover, the verbal simplicity of *SC* hides substantial complexity in its application since it has long been recognized that the use of a common term may involve a subtle ambiguity. It is only after such ambiguities have been eliminated that *SC* comes into play. In effect, there is a simplicity in the statement of *SC* that is not completely reflected in its application; in the case of *CV* this lack of simplicity is included in the proposition. But this similarity does not make it unimportant which of the two guiding assumptions we adopt. We should prefer *SC* if we believe that conceptual diversity is rare and that there is an underlying body of concepts that is shared by the members of conceptual communities. We should prefer *CV* if we think that there is substantial diversity among the concepts deployed by different individuals and that a significant portion of the concepts that bind conceptual communities together exhibit only a family resemblance.

*The Prevalence of Conceptual Diversity in Human Life*

I began my defense of *CV* by responding to objections, and in the course of this response I have also provided some positive evidence for *CV.* I now want to put aside specific philosophical debates for a moment and argue that conceptual diversity is more common in human life than is often recognized in the philosophical literature. I will begin with another suggestion from Putnam.

Putnam (1975b) argues that there is a division of labor in a community's knowledge of the meaning of a term. Many of us have only a limited grasp of the difference between an elm and a beech, but there are experts who have more detailed knowledge of these trees, and this expert knowledge contributes to a vector that constitutes a term's meaning. Like Putnam, I know little about the differences between elms and beeches; my knowledge is so limited that I cannot provide a principled basis for distinguishing them. Since this inability will make it difficult to explore the difference between my situation and that of an expert, I am going to shift to a different example. Until quite recently, I considered conifers and deciduous trees to be mutually exclusive classes and I did have a principled basis for the distinction: Conifers have needles rather than leaves, remain green all year long, and reproduce by means of cones. Deciduous trees have leaves that they lose every fall, grow new leaves in the spring, and reproduce by means of seeds.[7] However, I recently discovered that my classification does not match that used by botanists since their classification scheme includes deciduous conifers such as the larch. In other words, the concepts I associated with the expressions "conifer" and "deciduous tree" are not the same as the concepts that botanists associate with these words. This discovery led to a small revision in my botanical concepts – although it is likely that my revised concepts of a conifer and a deciduous tree still do not match those of an expert botanist. I want to note three points about cases of this sort.

First, experts in a field often know a great deal more about characteristic features of a set of objects than do nonexperts. At least some of this knowledge will be incorporated into the experts' concepts. Indeed, one clear indication of conceptual diversity is a principled difference in the organization of reference classes. In this respect, the present example is analogous to that of pre- and post-Copernicans who classify celestial objects differently as a matter of principle. Second, given some differences between my initial concept of a conifer and an expert's concept, if we were to analyze our respective concepts we might well come up with different results supported by different intuitions, and this difference would not automatically show that at least one of us had failed to provide a correct analysis. Third, these different concepts are not wholly disjoint.

[7] A botanist may note many confusions in this description. These can be added to the specific confusion discussed in the main text.

The botanists and I agree that we are talking about trees, and the botanist could use this agreement as a starting point in teaching her concepts to me. There would be no need for her to point out that dogs and oaks are not conifers.[8]

Differences of this sort are rather common. They occur wherever we have the kinds of variation in knowledge that distinguish those who have expertise in a subject from those who have a less developed understanding. Such differences can be generated wherever occupational or avocational interests lead some people to explore a subject in greater depth than others do. To borrow another of Putnam's examples (1975b, 241), those who know the chemical composition of jade and distinguish the two different minerals that are commonly classified under this rubric are using a different concept when they consider jade than do those who lack this knowledge – although these concepts are sufficiently similar to allow for communication and discussion of these differences. In the same way, a student of physical science will classify ordinary glass as a liquid while most will classify it as a solid; and a chemist will classify diamond and graphite as different allotropes of carbon, while most folks will consider them utterly distinct. Many people will not even have the concept of an allotrope and will thus have to acquire a new concept before they can grasp the chemist's account.

Let us pursue this last example a bit further. Whether a chemist considers diamond and graphite to be "the same" will depend on context and on her purposes. In some chemical contexts, where physical structure is not important, diamond and graphite will be considered the same "stuff"; in other chemical contexts they will be considered different. The recently discovered buckminsterfullerines are a previously unsuspected allotrope of carbon and led to the introduction of a new concept. Still, our chemist is not likely to consider graphite and diamond to be the same when selecting a gift for a lover.[9] This last example brings us back to a point made above:

[8] With regard to meanings Putnam writes: "Suppose our hypothetical speaker points to a snowball and asks, 'is that a tiger?'. Clearly there isn't much point in talking tigers with *him.* Significant communication requires that people know something of what they are talking about" (1975b, 248). From the perspective of the current paper, the conclusion that there is not much point in talking tigers with this person may be premature – it may only require extra work. This will depend on what concepts the person has available, such as whether she has a concept of an animal or a mammal.

[9] Isomers are the analog of allotropes for molecules that involve more than one element; isomers have the same chemical formula, but different properties because the elements are arranged differently. Since isomers involve more than one element, considerable complexity is generated and there are several types of isomers. Many biologically important molecules come in pairs that are mirror images of each other, but only one of these will be a normal part of our biochemistry while the other will be biologically useless. Further, two different isomers of a molecule may occur in an organism and play different roles. For example, the processing of a photon in the retina involves breaking a carbon double bond in one isomer of the compound retinal (11-*cis* retinal), rotation of part of the molecule, and re-forming of the bond to produce a different isomer (all-*trans* retinal). This

A single individual may have two or more overlapping concepts of a specific kind of item and shift between them as the context changes.

One consequence of the above remarks is that increasing our knowledge about a type of item often involves changes in the relevant concepts.[10] I am not claiming that learning always involves conceptual change, and I have no systematic account of when conceptual change occurs. For present purposes it is sufficient that conceptual change occurs in some cases, and (once again) I take it as clear that conceptual change has occurred when new information about some domain leads us to reorganize reference classes.

The examples I have just been considering primarily concerned conceptual diversity in a community at a given time; now let us look more carefully at changes that are introduced into a community over time. Early in the nineteenth century Dalton took a characteristic specific weight to be one of the defining features of each element. This view became established and chemists assumed that all pure samples of an element have the same specific weight. Differences in specific weights among samples of an element were taken as evidence that the samples were not sufficiently pure, and highly skilled chemists employed a great deal of effort in attempting to purify such samples. Once the existence of isotopes became understood, this necessary condition was dropped and the chemist's concept of an element was changed in a way that made much of this previous research chemically irrelevant.[11]

It is important to note that there was nothing internally defective in earlier versions of the concept of an element; the need for conceptual change arose out of a search for concepts that are instantiated in the actual world. This search is a major project of science. Sometimes earlier stages in this search have yielded concepts that we still consider adequate; sometimes concepts that are considered adequate at one time are simply abandoned; sometimes

change in the form of retinal alters the shape of the protein rhodospin, to which the retinal is initially bonded, and this change of shape generates a flow of ions that initiates an electrical impulse in the optic nerve (Silberberg 598). This brief account requires the introduction of several concepts that most of us will not acquire during the normal process of enculturation.

[10] Those who hold that concepts are individuated by necessary and sufficient conditions may object to talk of a concept being changed; for present purposes they may read such remarks as indicating that one concept has been replaced by a different – but similar – concept.

[11] Soddy (50) summed up the situation this way: "There is something, surely, akin to if not transcending tragedy in the fate that has overtaken the life work of that distinguished galaxy of nineteenth-century chemists, rightly revered by their contemporaries as representing the crown and perfection of accurate scientific measurement. Their hard-won results, for the moment at least, appear as of as little interest and significance as the determination of the average weight of a collection of bottles, some of them full and some of them more or less empty."

there is a long process of slow conceptual change leading us from one concept to its successors; sometimes concepts are introduced that are strikingly new. Along with science, technological innovation also generates a need for conceptual innovation. Lakoff (83–84) discusses an important example: We have long recognized that a woman who plays a central role in rearing a child need not be the child's biological mother and we have distinguished between a biological and a social mother. The new reproductive technologies have created situations in which the woman who contributes half of a child's genetic endowment need not be the woman who gives birth to that child, so we must now distinguish between a genetic mother and a birth mother. Reflection on the history of beliefs about the roles of males and females in reproduction, and on the development of our understanding of the basic biology involved, will reinforce the point that drawing this new distinction requires a great deal of conceptual background that was not available in the recent past. Moreover, newly introduced concepts will not be adopted by all members of a community simultaneously, so that diachronic conceptual change also generates synchronic conceptual variations among members of a community.

Differences of the sort we have been considering do not occur only as a result of scientific and technical developments. Ancient Greeks and Romans had rich systems of concepts for describing various deities and their functions; these concepts are not in the conceptual repertoire of most contemporary Americans. In the contemporary world, professional carpenters and plumbers will classify many materials and tools differently than will those who lack their detailed knowledge. One kind of difference will involve cases in which an expert would include items in a class that a nonexpert would initially exclude. For example, many people familiar with wrenches would require that a wrench be completely rigid and would have to make a small revision in their concept on encountering a plumber's chain wrench.[12] In a similar way, someone unfamiliar with standard nomenclature for lumber might expect two-by-fours to have a cross-section that is two inches by four inches, but this is not the concept that is instantiated in lumber yards, where "two-by-fours" are one and seven-eighths by three and three-quarters inches.

As a final example recall that, historically, ordinary English did not include a nonderogatory term to describe an unmarried adult female, while the term "bachelor" played this role for men. This linguistic fact is a reflection of an underlying conceptual fact: The very concept of an unmarried adult female implied a negative evaluation. Changing views of the roles of men and women in society virtually demand a conceptual

[12] An adequate theory of concepts would explain why, in this case, similarity of function plays a dominant role and leads us to classify the chain wrench as a wrench, while we do not reclassify a chair as a ladder when we stand on the chair to change a light bulb. Discussion of such issues is beyond the scope of the present paper.

innovation. One way to achieve this end is by changing the concept associated with the term "bachelor" so that the phrase "bachelor woman" is not a contradiction. Moreover, the need to use this compound expression might mark a passing phase in social development so that we can easily imagine a future situation in which the term "bachelor" has lost any connection to a particular sex. In these circumstances, viewing "All bachelors are male" as expressing a conceptual truth would be a historical relic on a par with viewing "All atoms are indivisible" as expressing a conceptual truth.

*The Future of Research in Conceptual Analysis*

My next task in this paper is to consider the kinds of research that would remain in the realm of conceptual analysis given *CV.* First, the internal coherence of concepts will continue to be an important concern. Russell's discovery that the established concept of a set was inconsistent provides an enduring reminder that general adoption of a concept does not guarantee internal coherence. *CV* also points to the importance of a second task: development and clarification of concepts that are in accord with our overall scientific and philosophic aims, quite apart from any question of whether these concepts are currently found in our community. For example, Tooley maintains that causal relations imply nothing about temporal relations between cause and effect, and part of his argument for this thesis derives from a wider philosophical agenda: He holds a causal theory of time and thus requires an analysis of causality that does not include any temporal predicates so that we can then use this analysis in an account of temporal concepts.[13] Now suppose that the concept of causality currently found in our community does require that cause precede effect in time, and that we do not commonly think of time in causal terms. But suppose also that a strong case can be made for holding that the direction of causality does determine the direction of time. We would then have good reasons for replacing our current causal concepts – at least in those contexts in which these issues are relevant. Our example also suggests a third function for conceptual analysis: the development of alternative concepts – with the understanding that some of these may turn out to be just what we need as we come to understand more about the universe.

Examination of disputes over the analysis of knowledge generated by Gettier will further support the above proposal. Before Gettier came on the scene there was widespread agreement that knowledge is justified true belief. Gettier's counterexamples – and the many variations that soon

[13] Tooley also notes, "There are a number of causal concepts" (252), and that although some of these can be reduced to others, which ones we take as basic depends on our views on several other topics, such as the nature of causal laws and whether we think that causal facts supervene on noncausal facts.

followed – generated widespread agreement that the task of analyzing knowledge was not completed by providing instances of justified true belief that are, intuitively, not cases of knowledge. There have been two main kinds of response to the Gettier examples: According to some philosophers, these examples show that justified true belief is only necessary for knowledge, not sufficient, and that an additional condition is required. Others hold that there is nothing wrong with the traditional analysis because the claims in question are not justified, but we need further work on the analysis of justification. There is, however, widespread intuitive agreement that the Gettier cases are not instances of knowledge.

How should we interpret these developments? Focus first on the view that an additional condition is required. On the standard interpretation of conceptual analysis, philosophers went from agreement that they had properly analyzed a shared concept of knowledge to agreement that they had all been wrong in their analysis, along with disagreement on exactly how the analysis should be fixed. At this point, substantial disagreement appeared among those who claim to be analyzing the same concept. Moreover, as proposals are made, new counterinstances proposed, and new clauses added to cover these counterinstances, the analyses have become ever more complex and convoluted. As Pollock notes (181), it is not likely that the common concepts we acquire as we grow up in a society and learn a language have this kind of complexity.

A similar story can be told about attempts to reanalyze justification. All of the Gettier-type examples involve cases of nondeductive justification, and our lack of an adequate understanding of this case is not news. The Gettier examples led those who located the problem in the concept of justification to initiate new lines of research on the nature of nondeductive justification. I suggest that this research aims at building a more adequate concept, not at peering more accurately into a preexisting concept or providing a more precise description of how people at large carry out nondeductive justifications.

Consider another example. Attempts to provide analyses of material-object language in a sense-datum language were motivated by the twin beliefs that the semantics of sense-datum predicates is less problematic than that of material-object predicates, and that sense-datum statements are more certain than claims about material objects. If these analyses had been successful, we might have had good reasons for adopting the resulting accounts of material-object language irrespective of whether this account actually captures the concept of a material object generally found in the community. Evidence that a sense-datum analysis does not capture everyday concepts gives no more reason for abandoning the sense-datum program than would the discovery that Aristotelian physical concepts are a close reflection of ordinary ways of thinking about the physical world provide a reason for abandoning the accomplishments of physics since the

seventeenth century.[14] The fact that the concept of a sense-datum is a sophisticated concept not found in ordinary language does not undermine this project. Further, although the project failed, this failure advanced our understanding of a number of issues in semantics and epistemology. None of these issues have anything to do with the way nonphilosophical, nonscientific folks conceive of these matters.

In addition, there are many concepts of interest to philosophers and associated with words widely used in ordinary talk, but for which we should have little concern with their everyday versions. Examples include simultaneity, infinity, gene, heredity, and force. There may be no coherent concepts associated with the everyday use of these terms, and even if there are, we have versions of these concepts that have developed out of science, and good grounds for believing that it is these alternative versions that should concern us. Historically, these developments may have begun from everyday concepts, but these concepts have been altered and refined to the point that only their sophisticated successors are now of serious interest outside of historical and anthropological studies. The everyday concepts may still be used as a starting point for teaching the refined concepts, but to a large extent, these everyday concepts should be replaced for philosophical purposes.

Indeed, there is no reason why philosophers who are reflecting on science or art or society or the human condition or any of the subjects that engage the intellects of reflective people should not propose new concepts for thinking about these matters. Thus one outcome of the approach I am proposing is that philosophical reflection on concepts should be recognized as a significantly more creative endeavor than it would appear to be on the usual accounts. Strawson once labeled this project "revisionary metaphysics" (9) and much early analytic philosophy involved a radical rejection of these endeavors. I am urging that conceptual revision and innovation constitute a respectable and important part of philosophy. Indeed, a major theme of this paper is that analytic philosophers have often been engaged in this creative endeavor, but have misdescribed it under the influence of *SC*.[15] The point is well illustrated by a remark that Sellars makes after completing his account of how the genius Jones introduced the concepts of thoughts and sense impressions into a Rylean conceptual scheme. Sellars tells us that "the spirit of Jones is not yet dead" and he locates this spirit in sense-datum theorists who, "with a modesty

[14] See Holland et al. (206–11) for a discussion of Aristotelian elements in intuitive physics.

[15] Reflection on the literature provides significant evidence that the changes I am proposing have been quietly taking place over the past few decades. While the original guiding idea of analytic philosophy was that the analysis of concepts is the only legitimate task for philosophers, it is now almost commonplace for analytic philosophers to note the passing of this thesis and to describe analytic philosophy as a diverse subject bound together by a commitment to clear argumentation. See, for example, Baille (ix–x), Heil, Quinton.

forgivable in any but a philosopher, [confuse their] own creative enrichment of the framework of empirical knowledge, with an analysis of knowledge as it was" (1963, 195).

Explicit adoption of *CV* would thus produce a significant change in the rhetoric of analytic philosophy while encouraging developments that have already begun. This is an important change because the prevailing rhetoric has substantive consequences. To the extent that philosophers are not just analyzing concepts common in their community, the traditional rhetoric amounts to an incorrect metaphilosophy – and this metaphilosophy has consequences for the practice of philosophical analysis. One important consequence is that philosophers do not systematically check whether their own concepts are widely shared. Jackson addresses this point in his recent defense of conceptual analysis. Replying to the suggestion that philosophers should use opinion polls to determine the general understanding of specific concepts, Jackson writes:

> My answer is that I do – when it is necessary. Everyone who presents the Gettier cases to a class of students is doing their own bit of fieldwork, and we all know the answer they get in the vast majority of cases. But it is also true that often we know that our own case is typical and so can generalize from it to others. It was surely not a surprise to Gettier that so many people agreed about his cases. (37)

I want to note several points about this reply. First, a professor's assessment of whether students in his course share his reactions is hardly an example of good polling methodology. Second, in the Gettier case Jackson claims agreement in the "vast majority" of responses. What about the outliers? Are they unimportant? In Kuhn's account of normal science the vast majority of relevant cases are dealt with successfully by the prevailing paradigm; resistant anomalies provide a primary reason for eventually questioning guiding assumptions. As the passage illustrates, another consequence of espousing *SC* is the failure to take evidence of conceptual variation seriously. Third, even granting wide agreement in the Gettier case, other examples we have canvassed suggest that this case is not typical. Nevertheless, Jackson makes a double extrapolation from this example: He extrapolates from a restricted sample of students to a larger population, and from the Gettier case to other concepts. Both of these extrapolations are subject to empirical test. After all, how can we know that our own intuitions are typical unless we check? Goldman comments on this point:

> Philosophers sometimes seem to assume great uniformity in epistemic judgments. The assumption may stem from the fact that it is mostly the judgments of philosophers themselves that have been reported, and they are members of a fairly homogeneous subculture. A wider "pool" of subjects might reveal a much lower degree of uniformity. (160)

Another example from the recent literature will underline how acceptance of *SC* limits metaphilosophical reflection on the aims of analysis. There is a current debate about the correct account of the concept of logical consequence, a debate with implications for both the content and philosophy of logic. Some participants in this debate describe their aim as providing a clear formulation of a universal pre-analytic concept of logical consequence, and they consider this pre-analytic concept to play a normative role with respect to technical developments in logic. Etchemendy (6) represents this approach:

> Briefly put, my claim is that Tarski's analysis is wrong, that his account of logical truth and logical consequence does not capture, or even come close to capturing, any pretheoretic conception of the logical properties. The thrust of my argument is primarily at the conceptual level, but again the main impact is at the extensional. Applying the model-theoretic account of consequence, I claim, is no more reliable a technique for ferreting out the genuinely valid arguments of a language than is applying a purely syntactic definition. Neither technique is guaranteed to yield an extensionally correct specification of the language's consequence relation. Needless to say, this conclusion requires that we reassess the intuitive significance of Gödel's completeness theorem, as well as the import of the failure of analogous results when we move, for example, to second-order logic.[16]

Clearly *SC* is being assumed; but if *SC* is false, then it is unclear what standards are being used in Etchemendy's evaluation of technical results in logic.

Hintikka provides a contrasting approach. He proposes a "new and better basic logic to replace ordinary first-order logic. This received basic logic owes its existence to an unduly restrictive formulation of the formation rules for first-order logic probably motivated by an uncritical acceptance of the dogma of compositionality" (ix). Note especially that Hintikka is not seeking an analysis of ordinary concepts, but rather a logic that is adequate for the analysis of mathematics – actual mathematics, not some pretheoretic view of mathematics. Thus, in his view, the key step on the road to a better logic is to ask, "What can logic do for a mathematician? What *is* the role of logic in mathematics?" (1). As the above reference to "formation rules" indicates, the key inadequacy that Hintikka identifies in standard first-order logic is in its descriptive function: "Its resources do not suffice to characterize fully such crucial concepts as mathematical induction, well-ordering, finiteness, cardinality, power set, and so forth" (7).[17]

[16] Hanson (366, n. 2) offers another analysis of this pretheoretic concept that results in significant criticisms of Etchemendy's account. However, it is not clear that Hanson and Etchemendy agree on where to look for this pretheoretic understanding.

[17] Here is one example of this lack of expressive power, an example that plays a large

Hintikka describes his new logic as a "conservative extension of ordinary first-order logic" (73), but it has consequences for several important issues. For example, the much debated axiom of choice is a purely logical axiom, and excluded middle is not valid. But the most important point for our purposes is that Hintikka's criterion for a correct logic is quite different from Etchemendy's. In Hintikka's approach, the practice of researchers who have been persistently involved in the deductive enterprise provides the starting point for reflective philosophical analysis. A correct account of the consequence relation derives from a logic adequate to this practice – not from some vague, unformulated concept.

We can get further insight into what philosophical research on concepts will look like if *SC* is fully abandoned by recalling two views that were already in the literature in the 1950s. The more conservative view is found in Carnap's account of explication. For Carnap, an explication is not just an explicit formulation of an available concept. Instead: "The task of *explication* consists in transforming a given more or less inexact concept into an exact one or, rather, in replacing the first by the second" (3). Development of this replacement concept involves a tradeoff between four criteria: similarity to the concept being explicated, precision, fruitfulness in the sense of being useful for the formulation of universal statements, and simplicity. With regard to the first of these criteria Carnap writes, "The explicatum is to be *similar to the explicandum* in such a way that, in most cases in which the explicandum has so far been used, the explicatum can be used; however, close similarity is not required, and considerable differences are permitted" (7). In other words, philosophical reflection on a concept should do more than just provide clear formulations of existing concepts with all their imperfections; philosophical reflection seeks to improve our conceptual situation. I consider Carnapian explication to be a legitimate and important task for philosophy, but it is not the end of the matter.

Sellars offers a more radical vision. It is a major thesis of Sellars's epistemology that the appropriate concepts for thinking about the world are no more automatically available to us than are the correct propositions for

role in Hintikka's argument, and one that logicians have been aware of for some time (cf. Quine 89–91). Consider two formalizations of the sentence "Everybody has a unique habit," taking persons as our universe of discourse:

(1) $(\forall x)(\exists y)(\forall z)(\exists w)(Hyx\&Hwz\&\sim Hyz\& \sim Hwx)$
(2) $(\forall z)(\exists w)(\forall x)(\exists y)(Hyx\&Hwz\&\sim Hyz\&\sim Hwx)$

There is no reason to prefer one of these translations to the other, but they are not logically equivalent because in (1) the quantifier $(\exists y)$ is in the scope of $(\forall x)$ while $(\exists w)$ is in the scope of both $(\forall x)$ and $(\forall z)$; these dependencies are reversed in (2). Hintikka's "independence friendly" logic is constructed to avoid such unwanted quantifier dependencies.

describing the world. The development of these concepts is one major task of both scientific and philosophic research, and pursuit of this project requires that we formulate and evaluate alternative systems of concepts. Thus Sellars tells us that "it is characteristic of modern science to produce deliberately mutant conceptual structures with which to challenge the world" (1953, 337). And in a later paper, he says, "The motto of the age of science might well be: *Natural philosophers have hitherto sought to understand 'meanings'; the task is to change them*" (1958, 288). According to the approach that I am proposing, this task is not limited to natural philosophers.[18]

I will end this discussion by considering some additional objections to the view I have proposed. One objection derives from the thesis that conceptual analysis seeks a priori necessary truths that are intrinsically deeper than any results arrived at by empirical research. Thus the kinds of historical and other empirical conditions I have been urging are just irrelevant. Note, however, that no matter how much conceptual diversity we are prepared to contemplate, we can easily generate necessary truths if we index each analysis to a specific concept deployed by a specific individual at a given point in time. Such necessary truths seem strikingly trivial since we can generate them at will through arbitrary definitions of terms. If I stipulate that by "framis" I mean "a male university professor less than six feet tall," then the proposition "Every framis is male" expresses a necessary truth that can now be discovered by a priori reflection. The triviality of such necessary truths is a familiar doctrine of the empiricist tradition – along with the view that all necessary truths are of this sort. This doctrine is regularly ignored by many who consider themselves to be working in this tradition, but who still maintain that necessary truths are deeper than truths discovered by empirical procedures. The necessary truths discovered by conceptual analysis would be especially deep if the concepts being analyzed were, for some appropriately deep reason, universal. In other words, the view that truths arrived at by conceptual analysis are especially important goes hand in hand with adoption of *SC* plus the thesis that the relevant conceptual community is, at least, all of humanity. As soon as we localize concepts in smaller communities, then conceptual analysis becomes the study of concepts that have become fixed in those communities. Presumably, an understanding of the reasons why these concepts have become fixed will require empirical study. These remarks should focus our attention on a crucial issue: arguments for believing that there is a single conceptual system underlying all human cognition. Examination of these arguments will have to be postponed for another occasion, but I think that the exam-

[18] See Brown 1986, 1991 for further discussion of these Sellarsian themes and additional examples. The entire approach will be developed in greater detail in a book, *Conceptual Systems*, that is currently in preparation.

ples canvassed in this paper provide reasons for doubting that such arguments will succeed.

Consider another objection: that many of the examples discussed here are inappropriate. We should not expect all concepts to be universal, but only a small, especially important set.[19] It may, for example, be particularly inappropriate to argue from natural-kind concepts since these are fairly clear examples of concepts that must be developed over time. But it is not clear how we are to identify the basic concepts. Examples I have discussed include such favorite topics of conceptual analysis as knowledge, causality, time, truth, and logical consequence. If my arguments are correct, these will have to be excluded from the class of universal concepts. In the absence of an account of which concepts are basic, I can only offer a willingness to discuss other specific examples as they arise.

Finally, I will note a variation on the last objection from the empiricist tradition: There is a large set of concepts, derived from outer and inner experience, which provide the material out of which all other concepts are constructed. These are the basic concepts, they are determined by our biology, and they may be the universal concepts. Two replies are in order. First, it is central to the empiricist tradition that these basic concepts are not subject to analysis; conceptual analysis concerns only the nonbasic concepts. Second, even if some version of the distinction between basic and nonbasic concepts is correct, the point remains that a major part of our conceptual repertoire consists of nonbasic concepts and the story I have been telling can be rewritten in terms of the introduction, modification, and abandonment of nonbasic concepts. The assessment of which nonbasic concepts should be included in our conceptual repertoire remains a major part of the process of understanding the world, our institutions, and ourselves.

## Conclusion

So why do conceptual analysts disagree? Sometimes because they have not arrived at a correct analysis of a shared concept; but sometimes, more often I suspect, because they are not attempting to analyze the same concept. There are many reasons why different people associate different concepts with a single word, and we have considered some of these, but I have not attempted to give a comprehensive account of the sources of these differences. I have, however, been especially concerned to emphasize that in an important subset of cases philosophers are not analyzing concepts at all, but rather attempting to develop concepts that are needed to carry out a philosophical agenda and thus arrive at a deeper understanding of some subject. The attempt to understand the world we live in

[19] Recall Socrates' doubts, expressed at *Parmenides* 130b–131a, that there are forms corresponding to every term.

is part of the traditional subject matter of philosophy, and the construction of better concepts for thinking about aspects of that world is a central part of this endeavor.

*Department of Philosophy*
*Northern Illinois University*
*DeKalb, IL 60115–2854*
*USA*
*hibrown@niu.edu*

**References**

Antony, Louise. (1993). "Quine as Feminist: The Radical Import of Naturalized Epistemology." In *A Mind of One's Own*, edited by L. Antony and C. Witt. San Francisco: Westview Press.

Baille, James. (1997). "Preface." In *Contemporary Analytic Philosophy*. Upper Saddle River, N.J.: Prentice Hall.

Bishop, Michael. (1992). "The Possibility of Conceptual Clarity in Philosophy." *American Philosophical Quarterly*, 29, 267–77.

Brand, Miles. (1980). "Simultaneous Causation." In *Time and Cause*, edited by P. van Inwagen. Dordrecht: D. Reidel.

Brown, Harold. (1986). "Sellars, Concepts and Conceptual Change." *Synthese*, 68, 275–307.

———. (1991). "Epistemic Concepts." *Inquiry*, 34, 323–51.

———. (1996). "The Methodological Roles of Theory in Science." In *The Scientific Basis of Geomorphology*, edited by B. L. Rhoads and C. E. Thorn. Sussex: John Wiley and Sons.

Carnap, Rudolph. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.

Etchemendy, John. (1990). *The Concept of Logical Consequence*. Cambridge, Mass.: Harvard University Press.

Feyerabend, Paul. (1962). "Explanation, Reduction, and Empiricism." In *Minnesota Studies in the Philosophy of Science III*, edited by H. Feigl and G. Maxwell. Minneapolis: University of Minnesota Press.

———. (1975). *Against Method*. London: New Left Books.

Fodor, Jerry, and LePore, Ernest. (1992). *Holism: A Shopper's Guide*. Oxford: Blackwell.

Gettier, Edmund. (1963). "Is Justified True Belief Knowledge?" *Analysis*, 23, 121–23.

Goldman, Alvin. (1992). *Liaisons*. Cambridge, Mass.: MIT Press.

Graham, George, and Horgan, Terry. (1994). "Southern Fundamentalism and the End of Philosophy." In *Truth and Rationality*, edited by E. Villanueva. Atascadero, Calif.: Ridgeview.

Hanson, William. (1997). "The Concept of Logical Consequence." *The Philosophical Review*, 106, 365–409.

Heil, John. (1995). "Analytic Philosophy." In *The Cambridge Dictionary of Philosophy*, edited by R. Audi. Cambridge: Cambridge University Press.
Hintikka, Jaakko. (1998). *The Principles of Mathematics Revisited*. Cambridge: Cambridge University Press.
Holland, John, et al. (1986). *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, Mass.: MIT Press.
Hume, David. (1978). *A Treatise of Human Nature*. 2nd ed. Edited by L. A. Selby-Bigge, revised by P. H. Nidditch. Oxford: Oxford University Press.
Jackson, Frank. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
Kant, Immanuel. (1963). *Critique of Pure Reason*. Translated by N. Smith. New York: Macmillan.
Kim, Jaegwon. (1995). "Causation." In *The Cambridge Dictionary of Philosophy*, edited by R. Audi. Cambridge: Cambridge University Press.
Kitcher, Philip. (1978). "Theories, Theorists, and Theoretical Change." *Philosophical Review*, 87, 519–47.
Kline, David. (1980). "Are There Cases of Simultaneous Causation?" In *PSA 1980*, vol. 1, edited by P. D. Asquith and R. Giere. East Lansing, Mich.: The Philosophy of Science Association.
Kuhn, Thomas. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
Lakatos, Imre. (1970). "Falsification and the Methodology of Scientific Research Programmes." In *Criticism and the Growth of Knowledge*, edited by I. Lakatos and A. Musgrave. Cambridge: Cambridge University Press.
Lakoff, George. (1987). *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.
Laudan, Larry. (1977). *Progress and Its Problems*. Berkeley: University of California Press.
Laudan, Larry, et al. (1986). "Scientific Change: Philosophical Models and Historical Research." *Synthese*, 69, 141–223.
Lee, Jig-Chuen. (1988). "The Nontransitivity of Causation." *American Philosophical Quarterly*, 25, 87–94.
Mackie, J. L. (1980). *The Cement of the Universe: A Study of Causation*. Oxford: Oxford University Press.
Maxwell, Grover. (1962). "The Ontological Status of Theoretical Entities." In *Minnesota Studies in the Philosophy of Science III*, edited by H. Feigl and G. Maxwell. Minneapolis: University of Minnesota Press.
Menzies, Peter. (1989). "A Unified Account of Causal Relata." *Australasian Journal of Philosophy*, 67, 59–83.
Miller, Richard. (1987). *Fact and Method*. Princeton: Princeton University Press.

Pollock, John. (1986). *Contemporary Theories of Knowledge*. Totowa, N.J.: Rowman & Littlefield.
Putnam, Hilary. (1975a). "The Analytic and the Synthetic." In *Mind, Language and Reality: Philosophical Papers 2*. Cambridge: Cambridge University Press.
———. (1975b). "The Meaning of 'Meaning.'" In *Mind, Language and Reality: Philosophical Papers 2*. Cambridge: Cambridge University Press.
Quine, W. V. (1970). *Philosophy of Logic*. Englewood Cliffs, N.J.: Prentice Hall.
Quinton, Anthony. (1995). "Analytic Philosophy." In *The Oxford Companion to Philosophy*, edited by T. Honderich. Oxford: Oxford University Press.
Ramsey, William. (1992). "Prototypes and Conceptual Analysis." *Topoi*, 11, 59–70.
Salmon, Wesley. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
Scheffler, Israel. (1967). *Science and Subjectivity.* Indianapolis, Ind.: Bobbs-Merrill.
Sellars, Wilfrid. (1953). "Inference and Meaning." *Mind*, 62, 313–38.
———. (1958). "Counterfactuals, Dispositions, and the Causal Modalities." In *Minnesota Studies in the Philosophy of Science II*, edited by H. Feigl, M. Scriven, and G. Maxwell. Minneapolis: University of Minnesota Press.
———. (1963). *Science, Perception and Reality*. New York: Humanities Press.
Silberberg, Martin. (1996). *Chemistry.* St. Louis, Mo.: Mosby.
Soddy, Frederick. (1932). *The Interpretation of the Atom.* London: John Murray.
Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.
Suppes, Patrick. (1970). *A Probabilistic Theory of Causality*. Amsterdam: North Holland Publishing Company.
Taylor, Richard. (1966). *Action and Purpose.* Englewood Cliffs, N.J.: Prentice Hall.
Tooley, Michael. (1987). *Causation: A Realist Approach*. Oxford: Oxford University Press.
Toulmin, Stephen. (1961). *Foresight and Understanding*. New York: Harper & Row.
van Fraassen, Bas. (1980). *The Scientific Image.* Oxford: Oxford University Press.
von Wright, Georg. (1971). *Explanation and Understanding*. Ithaca, N.Y.: Cornell University Press.