CONCEPTS

Humans are makers and users of representations. Prominent among these are concepts, or
the representations we employ in higher thought. When we formulate plans, construct
theories, explanations, and narratives, and reason more generally, we are engaging in
processes that manipulate concepts. Indeed, having concepts is necessary even to form a
belief, desire, or intention. To believe that the earth revolves around the sun requires
possessing the concepts *earth* and *sun*; to intend to sign a mortgage contract requires the
concept *mortgage*; and to desire a gin and tonic requires the concept *gin*. Without
possession of those concepts, and the means of combining them, thinking these thoughts
would be impossible. Thus concepts have often been termed the building blocks of
thought.

Concepts and Categorization

Among psychologists, concepts are thought of as the mental representations that enable
individuals to categorize objects in one way or another. For example, having the concept
elephant enables one to recognize elephants as such, to sort elephants from other
creatures such as rhinoceroses and wooly mammoths, and perhaps to label elephants
verbally. Categorization may be thought of as making a mental judgment that this thing is
an F, or as thinking that all of these Fs are Gs, as when one categorizes elephants as
animals.

1

Psychologists who study concepts have focused on the structure of these representations, and how they are learned and processed. The classical theory of concepts assumed that concepts are like definitions: they consist of sets of necessary and sufficient conditions for belonging to a certain category. To have the concept *elephant*, *electron*, or *prime number* is to grasp such a definition. This view derives in part from the idea that understanding what it is to be a certain kind of thing involves grasping the essential features of that thing, which is what proper definitions are supposed to summarize.

The classical view fell out of favor in the wake of W. V. O. Quine's attack on the existence of analytic definitions, as well as many empirical studies suggesting that people do not represent most categories in terms of necessary and sufficient conditions. Instead, people appear to treat many categories as if they were organized around similarity to a prototype. Prototypes are representations of the statistically typical or "normal" member of a category. The typical raven is black, but there are also albino ravens; typical tables are four-legged but there are also three-legged designs; and so on. When we think of ravens and tables, we often think of them in terms of these prototypical features, not in terms of a definition. Whether we take something to be a member of a category depends on how similar it is to the prototypical category member, not whether it possesses the category definition. Similarity-based views of concepts can explain many facts about people's judgments, including the greater speed and accuracy with which people categorize prototypical category members, the existence of borderline cases, and the fact that certain inductive inferences are facilitated for more prototypical categories.

Another perspective on concepts says that they are organized around theories. The theory-based view (sometimes called the "Theory theory") says that concepts either are

or are embedded in theories of a certain domain. Theoretical concepts like *quark* get their meaning from the theories they are part of (e.g., the Standard Model of particle physics), and so on this view do ordinary concepts like *elm* and *computer*. A central part of these ordinary theories is that they are committed to psychological essentialism. Essentialism means that people tend to believe that categories possess hidden or unperceived essences that determine the kind of thing that they are and that normally cause or produce their superficially observable traits. In distinction from the classical view, however, it is not assumed that we know what these essences actually are. Essentialist beliefs seem to emerge early: children will judge that something is really a cat even if it is painted and scented to appear just like a skunk, and will assume that something is a skunk if it has skunk parents no matter whether it looks and acts like a skunk. Concepts belonging to many domains seem to go beyond simple perceivable prototypical features, contrary to what many similarity-based views would predict.

As this brief sketch indicates, the psychological role of concepts is a complex one, and explaining how we use them to categorize will require appealing to many types of processes, including abstraction, judging similarity, and causal-explanatory reasoning.

Concepts, Publicity, and Possession Conditions

Concepts are widely believed to be public, meaning that it is possible for many thinkers to possess the very same concept. This demand goes back at least to Gottlob Frege, who argued that thoughts must be the kinds of things that can be grasped by many individuals, and by the same individual at different times. Publicity is important because in offering

3

psychological explanations we appeal to the thoughts that people have: someone's behavior is explained by their having certain beliefs and desires, for instance. These explanations have a degree of generality built into them, in that anyone who thought the same would behave the same. So psychological explanations, insofar as they are intended to be general, assume that concepts and thoughts are public. In light of this, theories of concepts need to explain their possession conditions in a way that satisfies publicity.

Many theories of possession conditions are inferentialist. Inferentialism is the claim that possessing a concept is a matter of being able to draw the right sorts of inferences. For example, someone who possesses *squirrel* can infer from how something looks and acts that it is a squirrel, and infer from the fact that it is a squirrel to the fact that it is a certain kind of animal, that it caches and eats nuts, and so on. To have the concept *plus* one needs to be poised to infer that 2 plus 2 equals 4, as well as indefinitely many other arithmetic propositions. An inferentialist holds that what it is to have a concept is just to be poised to use it in making such inferences.

Inferentialists must tell us something about which inferences a person needs to be able to make in order to possess a concept. Here there are two possible positions: holism and localism. Holistic theories claim that possessing any concept necessarily requires being ready to make many inferences involving that concept, and thus possessing any concept requires possessing vastly many others. To have the concept *apple*, and thus to be able to think that something is an apple, one needs to understand, among other things, that apples are a certain kind of edible fruit, that fruits contains seeds and have a certain characteristic reproductive and life cycle, that they have a range of salient properties and

4

uses, and so on. On holistic views, even grasping such a seemingly simple concept requires having a vast web of interrelated concepts and beliefs.

However, since concepts are individuated by their place in this web, any change anywhere necessarily alters some or all of the concepts that one possesses. Holistic theories have difficulty satisfying publicity, since almost no individuals will share all of their network of beliefs, and hence no two individuals will share exactly the same concepts. A task for holistic theories, then, is to explain what it means for concepts to be similar enough to share across differing belief systems.

Localistic theories claim that possessing a concept requires being poised to draw some inferences, but not anywhere near as many as holists claim. Hence localism says that possessing any concept requires having some others, but not a large number. The paradigm examples here are well-defined concepts: having *bachelor* requires only having *unmarried male*, since that is all that a bachelor is; the concept *lamp* might be exhausted by saying that it is *designed to give off light*. Perhaps while *red* is not a definable concept, one could not possess it without also possessing *color*. These concepts are ones that can be had simply by being poised to draw a few crucial or central inferences.

Localistic theories impose more stability than do holistic ones, since not every change in a person's beliefs also changes all of their concepts. However, localistic theories are also committed to the existence of analytic or conceptually necessary connections: these are the inferential links that are part of the very meaning of a concept. Many philosophers are skeptical of the existence of any such connections. As noted above, Quine argued forcefully that there is no reliable criterion for when we have a genuine conceptually necessary connection, as opposed to merely a deeply held empirical

5

belief. The price of adopting localism is accepting the existence of strong analytic conceptual connections.

Atomistic theories take a different approach. They claim that concept possession is not based on the inferences one draws with a concept, but rather with what that concept picks out in the world. Concepts for atomists are fundamentally a kind of category detector. A detector is a device that goes off or indicates to a creature when something is present in the environment. Having the concept *cat* or *red* is just a matter of having a mental detector that reliably tracks, co-varies with, or otherwise carries information about the presence of cats or red things in the environment. Because these detectors can reliably inform a creature about the world around it, this approach is sometimes termed "informational semantics," and so atomism may be thought of as an information-based rather than an inferentialist approach. Informational views and inferentialist views differ on whether the fundamental role of concepts is to detect categories in the environment or to facilitate inferences concerning categories.

The ability to reliably detect a category does not presuppose the possession of any other concepts in particular, so atomists do not need to posit the existence of conceptually necessary connections. Moreover, a detector for cats or red things can exist even if its connection to other concepts and beliefs changes, so atomism does not face the stability problem that holism faces. However, it is unclear whether many concepts are best thought of as detectors. Concepts like *recession*, *art*, *truth*, *quark*, *good*, and *negative number* are ones that it is hard to imagine simply detecting conditions in the environment, since they pick out highly abstract and theoretical categories that are by their nature

difficult or impossible to detect. Informational theories may be correct for some concepts, then, but not all of them.

Innateness and Concepts

The normal adult has a rich repertoire of concepts for living things and material substances, actions and other occurrences, tools and artworks, kinship and other social groupings, scientific and technical categories, and the framework of ordinary life in general. Infants and young children, however, possess very few of these, and certainly none that depend on language, acculturation, and pedagogy. Two important questions in understanding conceptual development are: what is the initial conceptual endowment that children begin with, and what processes exist to enrich that initial endowment?

The first question, in particular, arises in debates about nativism. Nativism itself is a contested notion, and there is no general agreement on what it means to call a psychological characteristic innate. However, it can be useful to understand claims of innateness as meaning that a characteristic is acquired by a non-psychological process; in particular, that it is not a *learned* characteristic. If it is agreed that learning is at least one way to enrich our conceptual repertoire, the first question becomes: which concepts do we learn, and which do we acquire without needing to learn them?

A classic answer advanced by empiricists such as John Locke and David Hume is that our innate concepts are just our sensory concepts: those that pick out immediately perceivable categories such as particular colors, textures, shapes, tangible and audible qualities, and so on. These only require the right sort of experience to be "awakened" in

us; all of our other concepts are then complex logical constructions out of these sensory primitives. *Snowball*, for example, might be a complex representation partially consisting of *white*, *cold*, *round*, and *hard*. The empiricist ideal is to extend this model to all of our concepts, such that everything in the mind is ultimately analyzable in terms of concepts derived directly from sensory experience. On this view, there would be relatively few primitive concepts and many complex ones.

Rationalists such as René Descartes and Gottfried Leibniz argue that empiricists are wrong to suppose that most of our concepts can be analyzed in sensory terms. Experience with the world may be needed to give us these concepts, but they are not properly speaking learned from these experiences. Modern-day rationalists like Jerry Fodor (1981) have turned this into an argument for the impossibility of learning any concepts whatsoever. Fodor's argument runs as follows. Assume that to learn anything is to acquire evidence in favor of a certain proposition. To learn that grass is green, for instance, is to find evidence that confirms this generalization about the natural color of grass. Learning, then, requires the ability to entertain this proposition. To do this requires having certain concepts: the concepts *grass*, *green*, etc. But if this is so, we must already possess these concepts so that we can frame the proposition to be confirmed. Similarly, with the case of learning a new concept, there must be some hypothesis involving the concept that we are trying to confirm. But by the reasoning above, any such hypothesis must be represented by the learner; which is to say that the learner must already have the concept that she is trying to learn in order to frame the hypothesis in the first place. So there cannot be any such thing as learning a concept, since any such act of learning presupposes the concept to be learned. This leads to a kind of radical concept nativism.

The contemporary debate between defenders of empiricism and rationalism has two major axes. First, there is the issue of how much of our conceptual repertoire is "copied" from sensory experiences. Neo-empiricists argue, by appeal to psychological and neuroscientific evidence, that all of our concepts are ultimately grounded in the senses. If they are right, higher thought turns out to be a complicated form of perceptual simulation. Their opponents maintain that most, if not all, of our higher thought takes place in an amodal system that is independent of these lower-level perceptual systems. This remains a live question, in part since the notion of what should count as evidence that a psychological process is "perceptual" versus "amodal" is unclear and contested.

Second, there is the issue of whether we can learn any genuinely new concepts. Advocates of concept nativism think the very idea of such learning is incoherent. Their opponents hold that the notion of learning as hypothesis-confirmation is unnecessarily restrictive. They propose a variety of other models on which we can acquire new concepts in a way that is appropriately adaptive and responsive to experience, but which does not assume that we already have the concepts in question. These "bootstrapping" processes might treat concept learning as akin to learning a new skill, or to adding a new vocabulary item to a language. Neither of these obviously involves hypothesis-confirmation in any problematic way. It remains open whether learning a concept is relevantly similar to these processes, and whether they can avoid the radical nativist's argument.

Conceptual Analysis and Folk Concepts

Concepts have played a significant role in philosophical methodology. A central tenet of so-called "analytic philosophy" is that the proper task of the philosopher is to analyze important concepts such as *cause*, *action, person*, *knowledge*, *free will*, *truth*, *meaning*, *good*, *justice*, and so on. A common goal of philosophical analysis is the discovery of non-circular necessary and sufficient conditions for falling under a concept. Hence the traditional analysis of *knowledge* claims that knowing that *p* involves having a justified true belief that *p*. Analyses of this kind play an important role in philosophical debate.

Analyses of concepts are often tested against thought experiments: real or fictional scenarios in which certain properties that are part of the analysis are present, and one asks whether the concept being analyzed properly applies in that situation or not. If it does not, then this shows that the analysis is somehow defective, since something may have the properties picked out by the analysis but not have the concept being analyzed. For instance, if *knowledge* simply is *justified true belief*, nothing could fall under the latter and not under the former. But as Edmund Gettier showed, seemingly decisively, there are cases of justified true belief that are not also cases of knowledge, and so the traditional analysis must be defective.

In recent years, experimental philosophers have challenged this method of doing philosophy. They have argued that the judgments that are elicited by these thought experiments are much less stable than philosophers have normally assumed. For instance, some of them show significant cross-cultural variation between Asian and Western populations, while others may vary across gender, degree of philosophical experience, and so on. Many conceptual analysts have assumed that their own intuitive judgments are representative of "our" concepts in general. But if these judgments are highly variable

and unstable, this assumption is false. The "folk" may not think the way philosophers have assumed that they do, or the way that philosophers themselves do. If this is true, then concepts may not be as public as has often been thought; and if philosophers remain interested in analyzing concepts, they must decide which population's concepts are of interest, and sample them more systematically. Seen in this light, there may be little reason to treat the concepts and intuitions of philosophers as being of any special interest. One response would be to continue to pursue the project of conceptual analysis, but to broaden it to cover a more inclusive set of empirically gathered intuitive judgments. Alternatively, philosophers might adopt a conception of philosophy that does not see it as primarily being in the business of analyzing concepts by means of thought experiments.

Open Texture, Vagueness, and Contestedness

Many of our concepts are neither sharp nor precise. This imprecision comes in several forms. Friedrich Waismann argued that most terms in natural language are "open-textured," by which he meant that the rules or conditions that govern their application do not determinately cover every possible case. We might encounter things that have some of the characteristics associated with falling under a concept such as *cat*, but others that are so irregular or eccentric that we are unsure whether to count them as category members or not. (Imagine something that looks and acts like an ordinary cat, but which keeps growing to a gigantic size, or which appears and disappears seemingly at random, or which occasionally transforms into something dog-like.) While for ordinary cases questions of category membership seem settled, there are indefinitely many cases where

11

our concepts simply do not specify how we are to "go on." The notion of open texture was also famously invoked by the legal theorist H. L. A. Hart to describe the way in which laws necessarily fail to cover every possible future case that might fall under them (e.g., "No vehicles in the park" depends on what may count as a vehicle in the future—are Segways and hoverbikes included?).

Vagueness exists where there are borderline cases of category membership. Whether a person is tall or bald is a vague matter, since while there are clear endpoints where membership is determinate, there are also many cases where it is not. A person with no hairs is clearly bald, but many people are shy of a full head of hair but not obviously bald. We may want to say that there is no fact about whether they are bald or not, or that the judgment that they are bald does not have a truth value. Vagueness and open texture differ in subtle ways. Both involve indeterminacy of application, and hence involve borderline cases. Vagueness in a concept can be eliminated by adopting some precise standard that fixes the boundaries of its application on a particular occasion of use, whereas open textured concepts have indefinitely many possibilities for indeterminacy in their application, and these possibilities are unpredictable in advance. In other words, we usually know the dimensions or features that must be made more precise to resolve vagueness, but for open textured concepts we have no set list of qualities to make precise.

A different kind of demarcation problem is raised by what W. B. Gallie referred to as "essentially contested concepts." Essentially contested concepts possess five features: (1) they are used to praise activities or achievements; (2) these achievements have internally complex features; (3) these features can be weighted in several possible

ways to determine how the achievement is assessed; (4) the way these features can be modified or weighted is empirically sensitive and open-ended; and (5) the concept plays a dialogic role in which each user of the concept knows that their use will be contested and challenged by other concept-users.

Gallie's examples of essentially contested concepts include *art*, *democracy*, *social justice*, and *Christian life*. All of these denote a certain kind of complex, multidimensional activity that possesses the relevant open-ended empirical character, but one where the component dimensions may be assigned different weight by different people. Moreover, these concepts have both aggressive and defensive uses, with partisans of differing value schemes pitted against one another. Consider debates over whether Duchamp's readymades, Pollock's action paintings, or Damien Hirst's "The Physical Impossibility of Death in the Mind of Someone Living" (a tiger shark floating in a glass and steel cage full of formaldehyde) fall under *art*. In legal philosophy and the social sciences the notion of essential contestedness has also found wide application. Further widely discussed examples include concepts such as *rule of law*, *power*, *racism*, and *medicine*.

While essential contestedness is similar to open texture and vagueness, there are important differences. Like cases of open texture, essentially contested concepts do not have their empirical domain of application settled in advance. There is room for new exemplars to alter the extension of the category that they pick out. Thus John Cage's *4'33''* may change what counts as a composition. On the other hand, open textured concepts typically have a single, unitary, and widely agreed on core of application. Essentially contested concepts do not, or at least need not. Partisans of two notions of *art*

or *democracy* need not agree on the central exemplars of those concepts, since they may give the relevant evaluative features different weights.

In terms of the psychological theories of concepts discussed earlier, vagueness can be captured by many similarity-based theories of concepts, since these inherently allow for borderline cases. Essentially contested and open-textured concepts, on the other hand, have a complex internal structure reminiscent of the theory-based view. The essentially contested nature of many philosophical concepts may also explain the fact that traditional philosophical debates seem so peculiarly intractable.

Daniel A. Weiskopf

Georgia State University

See also Analytic/Synthetic Distinction; Experimental Philosophy; Holism (Philosophy of Language); Inferentialism; Intentionality; Rule-Following; Theory Theory

FURTHER READINGS

Alexander, J. and Weinberg, J. (2007). 'Analytic Epistemology and Experimental Philosophy', *Philosophy Compass*, Vol. 2, pp. 56-80.

Block, N. (1986). 'Advertisement for a Semantics for Psychology', *Midwest Studies in Philosophy*, Vol. 10, pp. 615-678.

Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.

Cowie, F. (1999). *What's Within*? *Nativism Reconsidered*. Oxford: Oxford University Press.

Devitt, M. (1995). *Coming to Our Senses: A Naturalistic Program for Semantic Localism*. Cambridge: Cambridge University Press.

Fodor, J. (1981). The Present Status of the Innateness Controversy. In *Representations*, Cambridge: MIT Press, pp. 257-316.

Fodor, J. (1998) *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.

Fodor, J. and Lepore, E. (1992). *Holism: A Shopper's Guide*. Cambridge, MA: Blackwell.

Frege, G. (1956). 'The Thought: A Logical Inquiry', *Mind*, Vol. 65, pp. 289-311.

Gallie, W. B. (1956). 'Essentially Contested Concepts', *Proceedings of the Aristotelian Society*, Vol. 56, pp. 167-198.

Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis*, Vol. 23, pp. 121-123.

Hart, H. L. A. (1997) *The Concept of Law*. 2nd ed. Oxford: Oxford University Press.

Keefe, R. and Smith, P. (Eds.). (1997). *Vagueness: A Reader*. Cambridge, MA: MIT Press.

Knobe, J. and Nichols, S. (Eds.) (2008). *Experimental Philosophy*. Oxford: Oxford University Press.

Margolis, E. and Laurence, S. (Eds.) (1999). *Concepts: Core Readings*. Cambridge, MA: MIT Press.

Murphy, G. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.

Nagel, J. (forthcoming). 'Intuitions and Experiments: A Defense of the Case Method in Epistemology', *Philosophy and Phenomenological Research*.

Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA: MIT Press.

Prinz, J. (2002). *Furnishing the Mind*. Cambridge, MA: MIT Press.

Quine, W. V. O. (1953). 'Two Dogmas of Empiricism' in *From a Logical Point of View*. Cambridge, MA: Harvard University Press, pp. 20-46.

Waldron, J. (2002). 'Is the Rule of Law an Essentially Contested Concept (In Florida)?', *Law and Philosophy*, Vol. 21, pp. 137-164.

Waismann, F. (1951). 'Verifiability' in A. Flew (Ed.), *Logic and Language*. New York: Doubleday, pp. 122-151.