

Neoclassical Concepts

DEREK LEBEN

Abstract: Linguistic theories of lexical semantics support a Neoclassical Theory of concepts, where entities like CAUSE, STATE, and MANNER serve as necessary conditions for the possession of individual event concepts. Not all concepts have a neoclassical structure, and whether or not words participate in regular linguistic patterns such as verbal alternations will be proposed as a probe for identifying whether their corresponding concepts do indeed have such structure. I show how the Neoclassical Theory supplements existing theories of concepts and supports a version of analyticity and conceptual analysis.

After engaging in debate for a sufficient amount of time, one realizes that arguments can only proceed when participants share the relevant concepts. Establishing shared concepts is so important to philosophy that during the 20th century many philosophers explicitly stated that the entire aim of the field was to analyze and clarify concepts, helping to facilitate some debates and put an end to others. This article will argue that linguistic theories of lexical semantics have interesting and important consequences for the study of concepts. The first section sets out the phenomena a theory of concepts is supposed to explain, and the second provides a survey of some current theories and the challenges they face. Section 3 presents an overview of decompositional approaches to lexical semantics, and Section 4 shows how these approaches support a Neoclassical Theory of concepts that helps explain the phenomena at issue. Finally, Sections 5 and 6 develop a kind of conceptual analysis from this theory and show how it can be applied to philosophical problems.

1. Why Do We Need Concepts?

‘Concepts’, writes psychologist Gregory Murphy (2002, p. 1), ‘are the glue that holds our mental world together’. Murphy elaborates by describing how new things often exemplify a category that we are familiar with, so concepts act as a glue by connecting our ‘past experiences to our present interactions with the world’. However, ‘connecting our past experience to present interactions’ is vague, and it is appropriate to ask for a specific list of phenomena that a theory of concepts is supposed to explain.

As suggested above, concepts help to provide the common ground that all speakers share in a discussion. In a recent review, Rey (2010) writes: ‘Concepts are what

I wish to thank Steven Gross, John Waterman, Jonathon Hricko, and Jeff Maynes for helpful comments and feedback on earlier drafts.

Address for correspondence: Derek Leben, Department of Philosophy, University of Pittsburgh at Johnstown, 450 Schoolhouse Rd, Johnstown, PA 15904, USA.

Email: leben@pitt.edu

remain stable across epistemic variability, and so give argument a point'. Concepts are often described as the 'constituents of thoughts'; they are whatever enables my thoughts about elephants to be the same as your thoughts about elephants, and my wife's thoughts about elephants. Whether that is having a picture of elephants in our heads or being in the same sort of relation to actual elephants is a theoretical matter discussed in the next section, but for now we will motivate the need for concepts: why think there is some sort of stable information that relates all of our elephant-thoughts at all?

One of the primary phenomena at issue for theories of concepts is why there is such a close overlap between the information used when reasoning, categorizing, and drawing analogies about objects and events. For example, when making inductive inferences about what elephants are likely to do, one is likely to draw upon the same body of information (i.e. they are large grey animals with tusks) as when one makes comparisons between elephants and other objects (analogy) or groups perceptual cues under a single category by saying something like, 'ah, this is an elephant' (categorization). Concepts are often introduced to explain this relationship between the information in each of these capacities. Inference, categorization, and analogy-formation are often called 'higher cognitive capacities' to contrast with 'lower cognitive capacities' such as perception, motor skills, and proprioception. Categorization has historically been taken to be the primary capacity in which concepts are employed.

In order to support each of these cognitive capacities, individuals must have access to concepts. Fodor (1998) takes this to mean that concepts are 'mental particulars', while Machery (2009) makes a much more specific demand that concepts must be stored in long-term memory (Fodor's interpretation is non-committal about where and how concepts are stored). If they are mental particulars, then concepts must be acquired by some sort of learning or be part of our innate biological endowment. Finally, being mental particulars, there is presumably some sort of neural implementation of concepts, whether distributed or localized.

Putting together the phenomena that concepts are intended to explain, Machery (2009, p. 12) presents a sophisticated descriptive definition of what most psychologists commonly take concepts to be:

A concept of *x* is a body of knowledge about *x* that is stored in long-term memory and that is used by default in the processes underlying most, if not all, higher cognitive competences when these processes result in judgments about *x*.

I have been using the term 'information' as opposed to 'knowledge' to avoid any epistemic readings (following Machery's use of this term for the same purpose). In other words, when Machery says 'body of knowledge', he is not talking about justified true beliefs. The term 'default' here is supposed to be a theoretical term from artificial intelligence and computer science, and is intended to describe defeasible inferences, described as 'inferences that are normally drawn, except when some specific additional information is provided' (Machery, 2009, p. 11) It is often claimed

that concepts involve only a restricted set of the total information that an individual has about something. For instance, although a person may have access to a large amount of information about elephants, only a specific set of this is drawn on for use in higher cognitive capacities. Sometimes this is called a distinction between the mental 'dictionary' and the mental 'encyclopedia' (Komatsu, 1992, pp. 520–1).

In addition to accounting for the overlap in information used by higher cognitive capacities, concepts are also a plausible source of linguistic meaning. Sometimes linguistic meanings are simply called 'lexical concepts' to distinguish them from other concepts that happen not to be lexicalized (paired with syntactic and phonological structure). If linguistic meanings are just lexicalized concepts, then theories of concepts are a potential source for explaining aspects of linguistic meaning, such as compositionality and adicity.

Compositionality is the property of sentences such that 'the meaning of any complex expression is determined by the meanings of its parts and the way they are put together' (Pagin and Westerståhl, 2010). There are two observations about natural language that are taken to be evidence for compositionality: productivity and systematicity. Productivity is the principle that speakers are able to understand and produce a potentially unlimited number of (often novel) expressions, described here by Zoltán Szabó (2007):

Since competent speakers can understand a complex expression *e* they never encountered before, it must be that they (perhaps tacitly) know something on the basis of which they can figure out, without any additional information, what *e* means. If this is so, something they already know must determine what *e* means. And this knowledge cannot plausibly be anything but knowledge of the structure of *e* and knowledge of the meanings of the simple constituents of *e*.

For example, an English speaker can easily understand the meaning of an expression like: 'Alicia sent Christina the letter', even if that speaker has never heard it before. The explanation, presumably, is that the speaker understands the meaning of the components 'Alicia', 'sent', 'letter', etc., and the syntax of English, and can therefore understand the meaning of the larger expression (this works well for more exotic phrases).

Systematicity is a more controversial feature, where if speakers understand one arrangement of items, they can understand other arrangements of the same items 'for free'. Robert Cummins (1996, p. 594) provides a more formal definition:

A system is said to exhibit systematicity if, whenever it can process a sentence *S*, it can process systematic variants of *S*, where systematic variation is understood in terms of permuting constituents or (more strongly) substituting constituents of the same grammatical category.

Fodor's standard example is that if one understands the meaning of 'John loves Mary', one also understands the meaning of 'Mary loves John'. However, Johnson (2004)

has to my mind convincingly questioned whether this is a widespread phenomenon across language, since there appears to be a massive corpus of counter-examples from lexical semantics. Kristine can ‘fill the glass with water’ but not ‘spray the glass with water’. John can ‘put his gear down’ but not ‘stow his gear down’. Alice can ‘show Martha the book’ but not ‘describe Martha the book’.¹ Thus, when I discuss compositionality it will be exclusively based on considerations of productivity.

Another feature of linguistic meanings that they may inherit from concepts is adicity. Some linguistic meanings appear to be relatively ‘independent’, like words for objects (‘knife’, ‘chair’) and weather events (‘raining’, ‘snowing’). These words may appear on their own without much confusion or implication of any other words. Other meanings are what we might call ‘unsaturated’ (to use a metaphor from Frege), where they require the specification of other information. These include most words for events like ‘give’, ‘send’, ‘boil’, which all require that other information be included, like a giver and givee, a sender, recipient, and something sold, and something that is boiled. If a speaker utters the phrase ‘John sent’, there is some sort of information needed to complete or saturate this phrase for it to make sense. We will see later that this issue is complicated by the fact that many verbs describe apparently unsaturated events (like eating, which requires something eaten) but can be used intransitively (‘Bill ate’). Linguists have often offloaded the source of adicity to psychologists, assuming its source is in the concepts GIVE, SEND, and BOIL. Therefore, one might hope that a theory of concepts would explain exactly why some verbs involve one argument (participant), while others require two, three, or even four.

2. Theories of Concepts

The predominant view of concepts up to the middle of the twentieth century has been called the ‘classical view’ (Smith and Medin, 1981; Laurence and Margolis, 1999; Murphy, 2002), or the ‘definitional view’ (Fodor, 1998). Laurence and Margolis provide the following definition:

Most concepts—especially lexical concepts—are structured mental representations that encode a set of necessary and sufficient conditions for their application, if possible, in sensory or perceptual terms (1999, p. 9).

Murphy (2002, p. 15) observes how this view implies that ‘every object is either in or not in the category, with no in-between cases’, and all members of a category

¹ Johnson considers the response that these verbs actually belong to different grammatical categories, but he replies that this would make any formulation of systematicity tautological and completely uninteresting. I take this to mean that natural languages may be systematic, but the typical claim about the systematicity of language is a promissory note rather than a piece of empirical data to be explained.

are equally part of that category (i.e. there cannot be better and worse instances of a category). Alleged proponents of the classical view of concepts include John Locke and David Hume, who defined concepts like SUN or GOLD in terms of simpler perceptual features like ‘yellow’, ‘hard’, ‘round’, and so on. SUN and GOLD may share some necessary conditions like ‘yellow’, but they differ in sufficient conditions. Whatever these sufficient conditions may be, they would need to be accessible to everyone possessing the concepts. For instance, SUN would need to be defined in terms of features like: ‘large’, ‘bright’, and some sort of connection to being in the sky, as opposed to features like: ‘generates energy by nuclear fusion of hydrogen into helium’.

The classical view enjoyed its historical domination not only because of its inherent plausibility but because it provides an elegant explanation of each of the phenomena to be explained. To categorize a new item, one needs only to check all the perceptual features which constitute one’s concept against the perceptual features of that item. Inferences made on this structure can deductively analyze what features are ‘contained’ within the concept, or inductively build from these collections of features to a new perceptual prediction. Complex concepts are produced and acquired by combining simple perceptual features together in an iterative procedure. However, during the middle of the twentieth century the classical view was largely abandoned due to criticisms from both philosophers and psychologists that concepts do not have necessary and sufficient conditions provided by perceptual features (Wittgenstein, 1953; Quine, 1951; Rosch, 1975; Putnam, 1975; McCloskey and Glucksberg, 1978).

After the abandonment of the classical theory of concepts, the past fifty years have seen several replacements proposed. Following the discovery of typicality and gradedness effects in the 1970s, Rosch proposed that concepts are actually a statistical weighting of features called a ‘prototype’ (Rosch, 1975; Hampton, 1979; Lakoff, 1987). In the Prototype Theory, the concept VEHICLE would be a statistical weighting that encodes features like ‘carries things’, ‘moving’ ‘wheels’, etc. While some of these features may be perceptual, there is no reason why other features couldn’t be abstract (Hampton, 1998). For instance, Bloom (2004) proposes that the concept CHAIR may not only encode typical physical features like having a back and having legs, but also intended uses like being for sitting.

Psychologists endorsing the Knowledge Theory of concepts² (Carey, 1985; Murphy, 1985; Keil, 1989) agree that an open-ended list of features is constitutive of concepts, but instead of statistical relationships between features, the Knowledge Theory proposes that the relationships are explanatory ones, much like in scientific

² Again, this is to be read as ‘information’, not justified true belief. The Knowledge Theory is sometimes described as the ‘Theory-Theory’ of concepts, but I will avoid this term to avoid confusion with another theory in philosophy of mind with the same title, and for aesthetic purposes.

theories. Most Knowledge Theory proponents also abandon the distinction between any mental dictionary and mental encyclopedia.

Neo-Empiricist theories (Barsalou, 2008; Prinz, 2002) reject the assumption that there must be any relationship between the features of conceptual structure at all. Instead, they propose that collections of features are temporary ad-hoc constructions which are assembled from sensorimotor information to fit a given situation. One context may call for information about the color of elephants, while another may call for information about their size. In the first situation, the concept ELEPHANT would include color information, and in the second it would include size information, but there is no such thing as information that is a part of the concept in general (since there is no such thing as a concept in general).

While each theory has strengths, there are serious concerns about the ability of Prototype, Knowledge, and Neo-Empiricist theories to successfully explain all the phenomena described in the previous section. Just like many concepts do not have necessary and sufficient conditions, many concepts do not appear to have prototypes or perceptual features (e.g. PROTON, JUSTICE). They also cannot account for Putnam-style intuitions of concept possession without associated perceptual, prototype, or explanatory information (Putnam, 1962, 1975). Finally, it has been a common criticism that prototypes cannot account for compositionality or express adicity (Fodor and Lepore, 2002). This is because prototypes are not discrete entities but statistical weights between features. Fodor and Lepore allege that when combined, prototypes would simply blend together to create a new prototype rather than discretely combine in a way that preserves the structure of each.

In response to these concerns, Externalist theories of concepts have been proposed, where concepts have content about natural kinds external to them, and this external relation is itself constitutive of many concepts. Externalists may disagree about exactly how this relation comes about, whether the natural kind simply causes a mental particular x (Dretske, 1983), mental particular x was naturally selected in relation to some natural kind (Millikan, 1984), or there is some causal relation between mental particular x and natural kind y , with all other tokenings of x being metaphysically dependent on this causal relation (Fodor, 1987).

If one assumes a basic essentialism about kinds, Externalist theories do well in avoiding the challenges described above. The same information is drawn upon by higher cognitive capacities simply because there is a set of fixed information provided by that natural kind. Compositionality comes easily because presumably the natural kinds which concepts represent are discrete, and so composing concepts that are constituted by their relations to these kinds will also be discrete. We do not have to worry about the concepts HOUSE and FISH blending together, because actual houses and fish do not blend together. Most importantly for Fodor, intentionality is explained, since being related to actual fish is exactly what it is to possess the concept FISH. However, externalist theories have difficulty with other criteria for a theory of concepts. It is not clear why simply *being* in a specific relationship to elephants could help explain how individuals draw on the same information when reasoning,

categorizing, or making analogies about elephants. Externalists usually respond to this in the following way:

Perhaps the main thing that an [externalist] can say is that for any given concept, as much structure as you like can be invoked to explain its deployment [use in higher cognitive capacities], but with one serious qualification: The structure is to be treated as being merely associated with the concept rather than constituting a part of its nature (Laurence and Margolis, 1999).

Thus, Putnam (1975) maintains that in addition to a causal-historical connection between an utterance of the word 'tiger' and tigers, people also have a 'stereotype' (a structure similar to a prototype) that helps them draw inferences and make categorizations. Fodor (1975) appeals to 'meaning postulates' to do this work. Yet, if Externalist theories fail to account for some of the primary phenomena that a theory of concepts is indeed posited to explain, it is unclear why the stereotype and meaning postulates would not be proper components of concepts.

With Prototype, Knowledge, Neo-Empiricist, and Externalist theories successfully explaining different desiderata for a theory of concepts, there appear to be three options. One option is to try to stretch one of these theories to explain all the phenomena. A second option is to adopt a hybrid theory of concepts, where internal psychological components and external relations jointly constitute a concept. A final option is to simply abandon a unified theory of concepts altogether. Machery (2009) pursues this eliminativist route, using an analogy to the study of memory, which has discovered several heterogeneous systems underlying what might have previously been called 'memory'. These systems and processes include working memory, procedural memory, episodic memory, sensory memory, and so on. However, while Machery is correct that 'memory' may not be unified as a theoretical term at one level of explanation, the conjunction of these processes may still collectively constitute a theory of human memory at another, functional level. Similarly, I take it that a concept can be a conjunction of interconnected psychological components and external relations, so we will pursue the hybrid model here as a working assumption. The rest of this article will adopt the hybrid approach, and propose that work in lexical semantics provides a key part to a hybrid theory of concepts.

3. Lexical Semantics

Linguistic theories attempt to explain distributional patterns of language use (identified by either corpus-based or intuition-based methods) by hypothesizing common underlying entities and rules in the minds and brains of speakers. The null hypothesis is always that these patterns are driven by memorization and general knowledge,³

³ For example, one standard linguistics textbook (Akmajian *et al.*, 2001) describes syntactic and phonological entities and rules as only being warranted after first exhausting memorization, general knowledge, and abstraction as possible explanations.

but there are many patterns which cannot be accounted for by the null hypothesis. In addition to positing syntactic and phonological structures, some patterns require entities and rules which are much more specific to smaller classes of words. Consider one influential pattern: ways in which arguments of a verb can be re-arranged in syntactically acceptable ways, called 'verbal alternations'. In sentences containing some English verbs, speakers accept switching around the determiner phrases (DP) and prepositional phrases (PP) contained in the verb phrase (VP). For instance, English speakers regularly use and accept sentences like the following:

- (1) a. Derek sprayed water on the roses (Content-locative)
 b. Derek sprayed the roses with water (Container-locative)

Speakers also accept this kind of alternation with verbs from the following set: 'brush', 'cram', 'crowd', 'dab', 'daub', 'drape', 'drizzle', 'dust', 'hang', 'heap', 'inject', 'jam', 'load' (unless otherwise noted, all examples in this section are from: Levin, 1993; Pinker, 2007). If we limit ourselves to these verbs, one might assume that it is simply a general rule of English syntax that DPs and PPs within a VP are able to switch like in (1) above. However, if our sample is extended to other verbs in English, we quickly find that there are many verbs that do not allow this kind of alternation. Some verbs are capable of being used in the content-locative but not the container-locative:

- (2) a. Kristine poured water into the glass
 b. * Kristine poured the glass with water.

Obviously an English speaker would be able to infer what the speaker intends to communicate by an utterance of (2b), but there is something intuitively unacceptable about the sentence. This is not a single exception, but includes an entire class of words like the following: 'arrange', 'install', 'mount', 'put', 'hammer', 'pound', 'dribble', 'drip', 'pour', 'spew', 'coil', 'curl', 'twist', etc. Other verbs are capable of being used in the container-locative but not the content-locative:

- (3) a. * Kristine filled water into the glass.
 b. Kristine filled the glass with water.

Again, this is not a single exception, but includes an entire class of words like the following: 'blanket', 'block', 'coat', 'contaminate', 'cover', 'decorate', 'fill', 'frame', 'infect', 'litter', 'pollute', 'saturate', 'season', 'smother', 'surround', 'trim', 'veil', etc.

The initial hypothesis for this pattern is that speakers simply memorize which verbs are capable of entering into certain alternations and which verbs are not. This is certainly a possibility, but it fails because English speakers have reliable intuitions about which verbs will alternate and which will not, even when they have never heard these verbs used in that way. The second possibility is that speakers use 'general learning' to discover patterns in the verbs, extend these to novel verbs, and avoid constructions that they have not been presented with before. In the case of verb alternations, the patterns would need to be in the information that verbs are used to describe. For instance, the verb 'spray' is used to describe scenarios with two participants and a liquid. It could be the case that speakers learn patterns from

the information that verbs like ‘spray’ are used to describe, and identify this pattern with a certain construction like the locative alternation. The problem with this hypothesis is that verbs which are used to describe the same kinds of events can differ greatly in the kinds of alternations they participate in. For instance, ‘spray’ is used to describe an event that involves two participants and a liquid, but it appears that ‘pour’ and ‘fill’ can also describe events like this as well, without participating in the locative alternation.

Another illustration of this point is a puzzle involving the verbs ‘hit’, ‘cut’, ‘break’, and ‘touch’. Originally described by linguist Charles Fillmore (Fillmore, 1970), this has become a paradigmatic example of lexical semantics (Levin and Rappaport-Hovav, 2005; Pinker, 2007; Johnson, 2008). For our purposes, we can begin by noting that these four verbs all participate in the normal transitive construction (‘John hit/touched/broke/cut the bread’), and the events described by all four verbs are all very similar. They can be used to describe an actor moving, coming in contact with something, and affecting that thing. However, they participate in very different kinds of verbal alternations, labeled below:

- (4) Conative alternation:
 - a. Mary cut at the bread
 - b. Bill hit at the dog
 - c. * Nancy touched at the cat
 - d. * Jerry broke at the bread
- (5) Part-Whole alternation:
 - a. She cut him on the leg
 - b. She hit him on the leg
 - c. She touched him on the leg
 - d. * She broke him on the leg
- (6) Middle alternation⁴:
 - a. This bread cuts easily
 - b. * That wall hits easily
 - c. * This wire touches easily
 - d. This glass breaks easily
- (7) Causative alternation:
 - a. * The bread cut
 - b. * Mary hit
 - c. The vase broke
 - d. * John touched

⁴ There may be readings of the following sentences where they are grammatical, so we should clarify the middle alternation. (6a–d) should be read as describing properties of bread, walls, wires, and glasses, rather than actions of bread, walls, wires, and glasses. Thus, while ‘This wire touches easily’ might make sense when read as an action of the wire, it does not make sense when read as a property of the wire.

If speakers were simply memorizing verbs that fell into each alternation, we would expect there to be a fairly random distribution in a collection like: ‘break’, ‘hit’, ‘touch’, and ‘hit’. Although this may look random when isolated to these examples, there are many classes of verbs that act the same way. As Johnson (2008, p. 121) observes:

Verbs that pattern like *cut* in these constructions include: *hack, saw, scratch, and slash*. Patterning like *break* are *crack, rip, shatter, and snap*; like *touch* are: *pat, stroke, and tickle*; like *hit* are: *bash, kick, pound, tap, and whack*.

Further, if speakers were using patterns in the events described by the verbs to group them into alternating classes, we might expect that ‘break’, ‘hit’, ‘touch’, and ‘cut’ would all participate in the same alternations. Neither of these predictions is true.

After the null hypotheses have been defeated, theories of lexical semantics (like theories of phonology and syntax) posit hypothetical structures and formal rules in the mind and brain of speakers in order to explain linguistic patterns like verbal alternations. In this framework, each verb will have a specific structure that distinguishes it from other verbs, built out of the same set of basic elements. As described by Dirk Geeraerts (2010), there is a long history in linguistics to positing underlying semantic structures to words. The decompositional approach I will adopt is best exemplified in the work of Ray Jackendoff (2002), Steven Pinker (1989, 2007), and Beth Levin (1993; Levin and Rappaport-Hovav, 2005), and nicely summarized by Pinker (1989, p. 166):

Perhaps there is a set of semantic elements and relations that is much smaller than the set of cognitively available and culturally salient distinctions, and verb meanings are organized around them.

These elements, according to Levin and Rappaport Hovav (2005), are ‘isolated through an examination of the common semantic denominator of verbs exhibiting the same range of [distributions]’. In other words, lexical semantics takes classes of verbs that participate in a certain construction and looks for some common element, then repeats with other classes.

Semantic entities and rules should not be any more mysterious than the same entities and rules hypothesized in phonology and syntax. Phonological entities include ‘voiceless’ and syntactic entities include ‘auxiliary’ and ‘subject’. Lexical semantics posits entities like: AGENT, PATIENT, EVENT, STATE, THING, INSTRUMENT, CAUSE, PATH, PLACE, PROPERTY, and MANNER. For now, let us follow Pinker and call these entities ‘elements’. The term ‘semantic structure’ will be used to describe an organized combination of semantic elements. Semantic elements and structures may be realized in many ways, but whatever their ontology actually turns out to be, they exist in the same way that ‘auxiliary verb’ and ‘voiceless consonant’ exist in the mind and brain.

It becomes possible to explain verbal alternations in a simple and elegant way by appealing to underlying semantic structures. Taking the locative alternation, it was observed that there is no obvious information in the events described by these verbs that can set them apart, since events described by the verbs: ‘pour’, ‘fill’, ‘spray’ may all involve an actor, a liquid, and an affected entity. However, it is possible that the semantic structure of the verbs may focus on parts of what outwardly appears to be the same event. We may begin by asking: what common semantic denominator do verbs like ‘pour’, ‘vomit’, ‘nail’, and ‘twist’ all have in common? Lexical semantics suggests that it is an element we will call MANNER. This is inspired by the intuition that pouring a liquid onto a surface specifies the way in which an agent does an action, but not the way in which an entity is affected by that action (though obviously something is affected). This will have consequences for categorization; speakers agree that one can pour a liquid onto anything, but using the wrong manner of action (i.e. tripping and spilling some paint) is not correctly described by ‘pour’. How about verbs like ‘fill’, ‘blanket’, ‘cover’, ‘flood’, and ‘clutter?’ Decompositional theories propose that a word like ‘fill’ is paired with the concept FILL that is decomposable into (amongst other things) the entities GO and STATE, suggested by the intuition that they focus on a change in the state of the affected entity, but do not appear to care how that change of state is actually accomplished. Finally, verbs like ‘spray’, ‘load’, ‘spread’, and ‘cram’ participate in either construction because their semantic structures include all the elements: MANNER, GO, and STATE.

Figure 1 provides a summary of the hypothetical semantic structures posited in verbs to explain why they participate in some of the alternations described so far. There are also examples of some verbs that do and do not participate (indicated by *):

For an illustration of how this theory can be used to derive the semantic structure of a particular lexical item, consider the alternating behavior of ‘hit’ and ‘break’, which was so puzzling above. ‘Hit’ participates in the middle alternation, while ‘break’ does not. The opposite is the case for the part-possessor: ‘break’ participates in it while ‘hit’ does not. ‘Hit’ also participates in the conative, and ‘break’ participates in the causative. All of this suggests that the semantic structure of ‘hit’ contains the entities MOTION and CONTACT, while the structure of ‘break’ contains CAUSE, GO, and STATE.

Alternation	Semantic Elements	Examples
Locative	MANNER, GO, STATE	spray, load, *pour, *fill
Causative	CAUSE, INCH, GO, STATE	break, boil, *laugh, *cut
Conative	MOTION, CONTACT	hit, cut, *break, *touch
Part-whole	CONTACT	hit, cut, *break, touch
Middle	BECOME, GO, STATE	break, fill, *pour, *laugh

Figure 1 *The hypothetical semantic structures posited in verbs*

4. Lexical Semantics and Neoclassical Concepts

4.1 Conceptual Skeleton

I propose that theories of lexical semantics, which decompose word meanings into smaller components, support what will be called here a Neoclassical Theory of concepts. As the label suggests, this theory seeks to keep many benefits of the classical theory while introducing modifications that avoid its problems.

The first of these modifications is that the methodology of Locke and Hume involved introspection; to investigate the components of a concept like GOLD, one needs merely to reflect. Lexical semantics, as an empirical study of the mind and brain, does not rely on introspection but on hypotheses proposed to explain quantifiable linguistic and psychological effects. The structures it reveals are not necessarily accessible to conscious awareness (although they may turn out to be: see Johnson, 2004). Second, the Neoclassical approach does not demand that primitive elements of concepts be perceptual or drawn from experience, these can be entities like STATE and CAUSE which may very well be innate. Finally and most importantly, while the Classical theory proposes necessary and sufficient conditions for the possession of concepts, the Neoclassical view proposes only necessary conditions. Laurence and Margolis (1999, p. 54) propose the following definition of the Neoclassical view:

Most concepts (esp. lexical concepts) are structured mental representations that encode partial definitions, i.e. necessary conditions for their application.

As proponents and critics alike point out, this leaves us with a lot of ‘conceptual synonyms’. Take the verb ‘pour’, which has the same semantic structure as other verbs like: ‘drip’, ‘drizzle’, ‘spill’. In order to avoid the undesirable consequence that the concepts these words are paired with are all identical, there must be some kind of ‘completion’ in order to more fully distinguish these concepts. This completion is accomplished by the interface of elements in semantic structure with other information posited by other theories of concepts discussed above. The Neoclassical view is thus an important component of a larger hybrid theory, combining hypothesized semantic structures with other information from perception (Jackendoff, 2002), other cognitive systems (Pinker, 1989), or perhaps even person-external relations to objects and events (Rey, 1993).

The hybrid theory will therefore have components of Prototype, Neo-Empiricist, Knowledge, and Externalist theories, with semantic structures serving as a kind of ‘skeleton’ holding all of this information together. The metaphor of a skeleton is not new, and has been employed by Mohanan and Mohanan (1999), amongst others. In this metaphor, the completer information would be the ‘meat’, and various concepts might have different types of information fleshing out their lexical skeletons. Not all concepts have neoclassical structure, and the best probe for the existence of neoclassical structure is by looking at the distributional patterns of their lexicalized forms. As an example, I will focus on the concept CUT and its lexicalized form

(the verb ‘cut’), which has been a popular illustration in the lexical semantics literature. The verb ‘cut’ has a distinct pattern in the way it participates in alternations. It does participate in the conative, part-whole, and middle alternations, but does not participate in the causative. In addition to these alternation patterns, there are some other alternations that ‘cut’ participates in (Levin, 1993, pp. 156–7). I include some contrast cases as well to show some verbs which participate differently in these alternations:

- (8) Benefactive alternation:
 a. Derek cut a piece of turkey for Kristine / cut Kristine a piece of turkey
 b. The chef selected a meal for the critic / * selected the critic a meal
- (9) Instrument Subject Alternation:
 a. Tim cut the bread with a knife / The knife cut the bread
 b. Marty ate the ice cream with a spoon / * The spoon ate the ice cream
- (10) Characteristic Property of Instrument Alternation:
 a. This knife cut the bread / This knife doesn’t cut
 b. This hammer won’t break the window / * This hammer won’t break

The instrument subject alternation is similar to the causative alternation, except the instrument is moved to the subject position. The characteristic property of instrument alternation is similar to the middle alternation, but again, it is the instrument moved to the subject position. Both of these alternations suggest that information about an instrument is encoded in the structure underlying ‘cut’.

Because semantic structures are hypothesized to explain alternation patterns, we can infer from the patterns of ‘cut’ that the verb has the semantic elements: MOTION, CONTACT, INSTRUMENT, GO, STATE. Giving more detail about how to represent semantic structure requires committing to a specific formalism. While I wish to remain agnostic as to the proper formalization of semantic structure, I will follow Jackendoff (1990) in thinking of them as functions, and Pietroski’s (2005) proposal that they are organized by basic (truth-functional) conjunction. Thus, CUT (‘x cut y with z’) would be represented as the following:

- (11) CUT:
 AGENT(x) & PATIENT(y) & INSTRUMENT(z) & MOTION(x,z,
 linearly) & CONTACT(z,y) & GO(y) & STATE(y, separated)

Let us say (with Pietroski) that when predicates range over variables they are bound by existential quantification (not indicated here). Phrasing this in English is difficult, but one rough translation might be: *x* brings an instrument *z* into contact with *y*, changing the state of *y*. Some of this information is provided by context, like what instrument was used (a knife, a sword, a scalpel), who did the cutting, and what object was cut. Other information is more permanent information that completes the entities MOTION and STATE. In other words, exactly what kind of motion is involved? What state does the object go into? This information cannot be distinguished by alternation classes alone, but does fall within the scope of linguistics,

although here it does not seem like authors have much more to go on than their own intuitions. Hale and Keyser (1987) have speculated that the information paired with the STATE structure (which may be represented in the format of a prototype, sensorimotor pattern, etc.) is a ‘separation in material integrity’, and Pinker (1989) and Higginbotham (1989) propose that the MOTION is going ‘against and through’, and ‘linear[ly]’, respectively.⁵

4.2 Information Used by Higher Cognitive Capacities

When presented with a new item in experience, the classical view proposes that an individual can compare the features of that experience with perceptual conditions that constitute conceptual structure and discover a corresponding concept that matches this new experience. Similarly, the Neoclassical Theory proposes that one can (unconsciously) look for matches between the features of an experience and elements in semantic structure to narrow down a list of potential concepts.

The semantic structures hypothesized to explain linguistic and psychological effects provide a surprisingly effective method for categorizing novel items in one’s experience. For example, when presented with a cutting situation, the semantic structures: CONTACT, MOTION, INSTRUMENT, GO, and STATE will be cued. Of course, this would all be done unconsciously, as opposed to the classical view. The list of possible concepts would thus be narrowed down to CUT, SPRAY, and other members of their alternation class. At this point, the speaker may (again, unconsciously) use the ‘completer’ information (Material separation of the object? The instrument going against and through in a linear way?) to further determine whether that situation is a CUT situation or not. Despite this need for completer information, the fact remains that a significant amount of specification for categorization is already done by semantic structure. Pinker (1989, p. 108) provides some examples of categorization:

... causing an incision to appear by hard tugging or rapid heating followed by freezing is not cutting ... if one were to wave a knife in the air as part of a magic spell, causing the bread to split, that would also not be a clear example of cutting the bread.

Higginbotham (1989, p. 469) also provides some examples:

... one cannot be said to cut syrup at room temperature by passing a knife through it, because you cannot separate it by this means; you cannot cut a fish by punching a hole in it, or by stapling it to the wall, because the separation

⁵ Pinker (1989, p. 199): ‘Bob acts on a pear, causing the pear to become cut, by means of acting on a knife, causing the knife to go against and through the pear’. Higginbotham (1989, p. 4): “‘cut’ is a V that applies truly to situations e, involving a patient y and an agent x who, by means of some instrument z, effects in e a linear separation in the material integrity of y’.

in the fish's body is not linear; if fish had zippers you couldn't cut them by unzipping them (any more than you can cut a briefcase by just unzipping it), since the fish's material integrity would not be disturbed thereby; and so on.

John Searle (1980) makes a number of claims about how to categorize instances of CUT as well, but he bases these on what might be plausibly called a Knowledge Theory of concepts. Some of Searle's claims overlap with the Neoclassical Theory; for instance, he notes that stabbing a lawn with a knife or running over a cake with a lawnmower are *not* correct applications of CUT. Searle claims that this is because these situations do not fit with general knowledge about how people typically cut things. A Neoclassical Theory, on the other hand, checks these situations against the necessary features for CUT and finds these features missing. In stabbing the lawn, the state brought about is not separation, and in both stabbing and mowing the lawn, the motion is not moving against and through in a linear way.

Other cases of categorization provided by a Neoclassical theory inspired by lexical semantics are relatively easy to produce. One can fill a receptacle in just about any way, with just about anything (solids or liquids), because the semantic structure only constrains the way the concept FILL is applied to states of the receptacle. In contrast, one can pour liquids on just about anything: containers, surfaces, objects, even black holes. These all count as instances of the concept POUR, as long as they satisfy the features specified by the MANNER element of the verb. It doesn't matter how one sends something, so long as some object moves to a new state. Even when something is sent almost instantaneously, like an email, the concept still requires some (even only imagined) transit. If someone already possesses something, you cannot send or give it to them. One could say something like 'send me this email again', just like one could say 'send me this book again', but in both cases the interpretation seems to be to send a copy (another token of the same type) rather than the same exact object.

4.3 Properties of Linguistic Meaning

The most important thing that semantic structure contributes to lexical concepts is discreteness. A discrete system is one that uses discontinuous values, as opposed to an analogue system that employs continuous ones. The use of non-continuous elements is an important way of organizing information as seen clearly, for example, in the discrete elements of the genetic code. While features like prototype structures are analogue, semantic structure provides a discrete skeleton onto which prototype, sensorimotor, and knowledge information can be molded. The discreteness provided by semantic structure may explain properties of linguistic meaning like productivity and adicity. The problem of productivity was that all theories of concepts except the Externalist Theory fail to specify the boundaries of concepts in a way that prevents them from blending together (Fodor and Lepore, 2002). Yet this is exactly what the 'skeleton' of neoclassical structure allows. When words like 'send' and 'letter' are combined, the real combination takes place at the level of semantic structure, where 'letter' takes the position for THING present in the structure of SEND. Thus, only a subset of information associated with the concept

LETTER is specified: a set of information that is the ‘common denominator’ with other words that can take the same position (exhibit the same linguistic patterns). Other information associated with LETTER is not a part of this process, like the fact that letters are often lightweight.

Turning to adicity (or argument structure), a plausible theory of why some verbs require two participants while others require three is that the events, or concepts of the events, typically involve two and three participants, respectively. For example, why is it that giving and sending verbs (‘give’, ‘send’, ‘ship’) allow three subsequent arguments, touching verbs (‘touch’, ‘hit’, ‘stroke’) two, bodily motion verbs (‘smile’, ‘sneeze’) one, and weather verbs (‘rain’, ‘snow’) arguably none? Because the events of giving and sending involve three participants, touching involves two, bodily motion involves one, and weather none. This explains a wide variety of adicities as well as the plausible connection between argument structure and semantics.

However, there are problems with the simplistic mapping between adicity and events: namely, some verbs which appear to denote identical events can have different adicities. Common examples include:

- (12) a. John gave the museum the money
b. * John donated the museum the money
- (13) a. * Nick devoured
b. Nick ate
- (14) a. * The city destroyed
b. The city exploded

Therefore, the source of adicity still remains a problem.

A Neoclassical Theory of concepts can help solve this problem by appealing to the neoclassical structure of the verbs. While ‘give’ and ‘donate’ both appear to describe extensionally equivalent events, extensionally equivalent verbs can still have differences in neoclassical structure. By hypothesizing this level of structure, we can explain why ‘give’ has three participants: it specifies both GO and HAVE, while ‘donate’ focuses on MANNER. We posit these structures because ‘donate’ shares similar alternation patterns as verbs with MANNER, like ‘pour’:

- (15) a. Kristine poured water into the glass / donated money to the museum
b. (*) Kristine poured the glass with water / * donated the museum money⁶

Similarly, ‘give’ shares similar alternation patterns as verbs with GO and HAVE, such as ‘send’. One of these alternations is the dative alternation:

- (16) a. Sarah sent a letter to Neil / gave a letter to Neil (Prepositional dative)
b. Sarah sent Neil a letter / gave Neila letter (Double-object dative)

⁶ This sounds even worse if ‘the museum’ is replaced with a pronoun.

Decompositional theories propose that both ‘give’ and ‘send’ involve the semantic structures HAVE, GO, and STATE, inspired by the intuition that they can focus on either the part of the event that includes a person possessing something, or the way that the object travels. Compare this to other verbs which only participate in the prepositional dative but not the double-object dative like: ‘cry’, ‘groan’, ‘grumble’, ‘howl’, ‘purr’, ‘shout’, ‘stammer’, ‘whistle’, etc.:

- (17) a. Aarthi whispered the message to him
 b. *Aarthi whispered him a message

Lexical Semantics proposes that this is because the structure underlying these verbs only contains GO and MANNER.⁷ Other verbs are capable of being used in the double-object dative but not the prepositional dative because they can only focus on the possession itself (i.e. they only include HAVE):

- (18) a. *Nisha asked a question to Preetham
 b. Nisha asked Preetham a question

Verbs that act like ‘ask’ include: ‘cost’, ‘deny’, ‘envy’, ‘forbid’, ‘forgive’, ‘guarantee’, ‘issue’, ‘refuse’, ‘save’, and ‘wish’, etc.⁸ This account of adicity preserves the intuitive connections between verbs and conceptual structure, while adding an intermediary level that accounts for the mismatches and problem cases.

5. Conceptual Analysis

An important consequence of the classical theory is that when deductive inferences are applied to conceptual information, the resulting conclusions have a number of interesting features like being true by virtue of meaning alone, necessarily true, and knowable a priori. This is the origin of what has come to be called ‘analyticity’, but are what Locke (1690, IV, 8) originally called ‘trifling propositions’, because they ‘bring no increase to our knowledge’. Locke proposed that this taking apart of the concept occurs when ‘a part of the complex idea is predicated of the name of the whole’.⁹ For instance, for a concept like GOLD which is built up from perceptual features, making a claim like ‘Gold is yellow’ is simply taking apart a concept into

⁷ This proposal is supported by the syntactic similarities between the prepositional dative and the content locative, which are both hypothesized to involve MANNER.

⁸ This would explain why it sounds strange for inanimate objects to appear in the double-object dative construction, like # ‘Mary gave Chicago a letter’.

⁹ Kant (1781, A7, B11) appears to follow the Empiricists in his formulation of analyticity and the role it can play. While Kant formulates analyticity in several ways, the most famous (and consistent) is a straightforward containment relation similar to that of the Empiricists:

In all judgments in which the relation of a subject to the predicate is thought ... Either the predicate B belongs to the subject A as something that is (covertly) contained in this concept

its essential sensory components. We can think of this formally as and-elimination on one line to deduce a subsequent line:

Yellow & Malleable & Shiny & ...
 Yellow

During the late 19th and early 20th centuries, the analytic/synthetic distinction received a new formulation as well as a more important role. The new formulation was ushered in by Gottlob Frege (1884), who defined analytic inferences in terms of logical truth and synonymy. A logical truth is a statement or collection of statements where any non-logical expression can be freely interchanged with any other non-logical expression while preserving the truth-value of the statement. Rey (2008) provides the following examples of logical truths:

- (19) a. All doctors that specialize on eyes are doctors
 b. All cats that chase mice are cats.

Claims (19a-b) are logical truths because one can switch around the referring expressions like 'doctors' with 'cats', or 'eyes' with 'mice', and preserve the same truth value. The notion of synonymy comes in when considering expressions like (20):

- (20) All ophthalmologists are doctors

Sentence (20) is supposedly analytic because one can define ophthalmologist as a doctor that specializes on eyes, and this can be a logical truth when predicated of doctors. W.V.O. Quine (1951) delivered an important criticism to this formulation of analyticity as synonymy for conceptual truth, showing that there is no way to explain synonymy without appealing to necessity, which is one of the features that Frege's formulation of necessity set out to explain (thus rendering Frege's version of analyticity circular).

Quine's criticisms are widely thought to have destroyed the analytic/synthetic distinction, but they really only challenge the formulation of it developed by Frege. There have been some defenders of analyticity who reject this formulation in favor of a more psychologistic version along the lines originally proposed by Kant and the British Empiricists. As Horwich (1992) points out, Quine's assumptions about behaviorism and the publicity of language prevented him from being able to realistically address this early modern version of analyticity. Lyons (2005) suggests: 'If we just take concepts to be psychological entities, in particular, those mental representations

A; or B lies entirely outside the concept A, though to be sure it stands in connection with it. In the first case I call the judgment analytic, in the second synthetic.

Like the British Empiricists, Kant also appears to consider analytic inferences as trivial. He assumes that everyone will grant that analytic judgments are knowable a priori, and wants to go on to establish that aspects of synthetic judgments (claims about mathematics, the natural world, and morality) are also knowable a priori.

the tokening of which is partially constitutive of the tokening of an occurrent belief, the notion of containment becomes tractably clear'. Such a neoclassical approach to analyticity has been recently developed and discussed by Chomsky (2000), Horwich (1992), Lyons (2005), Pietroski (2003), and Rey (2008), which each author calls 'I-analyticity', inspired by Chomsky's distinction between E-Language and I-Language. I-analyticity holds speakers to have, often implicitly, certain basic inferential commitments by virtue of understanding their native language.

By taking a cognitive view of concepts and linguistic meanings it becomes possible to reconsider the traditional account of analyticity in terms of deductive inference on neoclassical concepts. For the present purposes, I am assuming that this capacity for deductive inference performs in a way that instantiates classical logic, and I will use normal notation and titles for deductively valid inferences like and-elimination and modus ponens. We have considered the neoclassical structure for CUT, which includes the elements: MOTION, CONTACT, GO, STATE. We also assumed this semantic structure is paired with sensorimotor, prototype, and knowledge information about the instrument going against and through the object, which is then in a state of being separated in material integrity. All of this information constitutes the concept CUT, and supports categorization, analogies, and combination with other conceptual information. I propose that the process of deductive inference can produce conclusions from information in neoclassical structure. For instance, from the neoclassical structure for CUT described above, any of the following conclusions could be potentially drawn:

- (21) a. An instrument touched the pear
 b. The pear is now separated
 c. An agent used the instrument (intentionally)
 (...)

To make the process of deductive inference from neoclassical structure explicit, we should consider the hypothetical steps involved at each stage. Let us say that a speaker hears an utterance of the sentence: 'Sean cut the pear', and then subsequently has the following semantic structure tokened in her mind and brain by filling in the semantic structure from (11), producing the following (once again omitting the existential quantification):

- (22) AGENT(Sean) \wedge PATIENT(pear) \wedge INSTRUMENT(z) \wedge MOTION
 (Sean, z) \wedge CONTACT(z, pear) \wedge LINEARLY(Sean, z) \wedge GO(pear) \wedge
 STATE(pear, separated)

Then, by the process of and-elimination on (22), another semantic structure is produced in the speaker's mind and brain:

- (23) INSTRUMENT(z) $\&$ CONTACT(z, pear)

Now it is possible to match this semantic structure with the appropriate syntactic and phonological information to produce an utterance like: 'An instrument touched the pear', or 'The pear came into contact with an instrument', etc.

An important point that should be emphasized is that deductive inferences can operate over premises containing a mix of neoclassical, completer, and non-conceptual information. Consider the sentence: ‘Sarah gave Neil a subpoena’, which is paired with the corresponding neoclassical structure:

(24) AGENT(Sarah) \wedge GO(subpoena, Neil) \wedge HAVE(Neil, subpoena)

Assume that the speaker has general knowledge about subpoenas, and is aware that just by possessing them one is legally required to appear before the court. Taking this general knowledge about subpoenas and combining it with the semantic structure in (24), it is possible to deduce as follows:

- (25) If Neil has a subpoena, he is legally required to appear before the court.
- (26) HAVE(Neil, subpoena) (24, and-elimination)
- (27) Neil has a subpoena.
- (28) Neil is legally required to appear before the court. (25, 27, modus ponens)

Although I-Analyticity in this case is analyticity of inferences (specifically deductive inferences over neoclassical structure), it can be easily expressed in sentences which lexicalize this inference (which would then be I-analytic), such as: ‘If Sarah gave Neil a subpoena, then Neil has a subpoena’. Presumably most real-world inferences are more like (24–28) than (22–23), combining both conceptual and factual information to produce interesting and valuable conclusions. Yet the types of inference that we are presently interested in are those rare instances where all the information in a deductive inference is drawn from neoclassical structure. These cases are remarkably similar to what Locke and Kant described when they spoke of inferences where the conclusion of an inference is somehow contained in the premises.

6. Baking Philosophical Bread

Putnam (1962) and even Quine (1951) conceded that there may be salvageable formulations of analyticity, but both questioned whether any of them can ‘wash philosophical windows or bake philosophical bread’ (Putnam’s expression). As Gillian Russell (2007) remarks: ‘In general, the weaker the conception of analyticity, the harder it is to argue that analyticity is philosophically interesting, and the stronger the conception of analyticity, the harder it is to argue that such a property really exists’. As far as conceptual analysis is concerned, even if we can discover the ‘skeleton’ of event concepts like CUT, BREAK, SPRAY, SEND, and TOUCH, there are not many philosophical debates centered around these concepts. Further, even if there are philosophically interesting concepts which enter into alternation patterns, the Neoclassical Theory of concepts is only a partial explanation of categorization and linguistic meanings. So can the Neoclassical Theory give us any information about philosophically interesting concepts, and what’s more, any *useful* information?

As Chomsky (2000) observes, lexical semantics has little to say about the concepts GOD, DEMOCRACY, or HAPPINESS, because there are no interesting

distributional effects that differentiate these words, preventing us from positing any semantic structure. However, concepts like KNOW and CAUSE may be productively analyzed. To some extent, philosophers working in analytic epistemology have been converging towards a Neoclassical Theory in their investigation of the semantics of the verb ‘know’. Contextualists like Cohen (2001) claim that ‘know’ is context-sensitive in the same way that ‘flat’ and ‘bald’ are. However, Jason Stanley (2004) argues that Contextualism is incorrect because ‘know’ does not belong to the same semantic categories as ‘flat’ and ‘bald’, as determined by patterns which cannot be explained by syntactic rules and which seem to specify information about the word’s content (in other words, the lexical semantics of ‘know’). Specifically, gradable terms take modifiers and allow for comparatives, while ‘know’ does not:

(29) Modifiers:

- a. * John very knows that penguins waddle.
- b. * John knows very much that penguins waddle.

(30) Comparatives:

- a. (?) Hannah knows more than John that Bush is president.
- b. (??) Hannah knows that Bush is president more than she knows that Clinton was president.

Stanley concludes that Contextualism is incorrect and ‘know’ is not gradable. Looking at what Stanley is actually doing here, he is using constructions which cannot be explained by syntactic rules to posit information in the semantic structure underlying ‘flat’ ‘bald’, and ‘know’. Very broadly, this is using a Neoclassical approach to do conceptual analysis.

Is there a way to use the kind of verbal alternation classes discussed in this article to address philosophical debates? Indeed there is; I think that some of the most important and difficult philosophical debates involve conceptual disagreements about causation, possession, and agency, and these debates can be usefully illuminated by studying the semantic structures underlying words used in these debates. Pinker (2007) recounts the story of the assassination of US President James Garfield:

On July 1, 1881, President James Garfield was waiting to board a train when Charles J. Guiteau took aim at him with a gun and shot him twice. Both bullets missed Garfield’s major organs and arteries, but one lodged in the flesh of his back. The wound was minor by today’s standards and needn’t have been fatal even in Garfield’s day. But his doctors subjected him to the harebrained medical practices of the time, like probing his wound with their unwashed hands (decades after antisepsis had been discovered) and feeding him through his rectum instead of his mouth.

At Guiteau’s trial, he reportedly remarked: ‘The doctors killed him, I just shot him’.

Ethicists are concerned about cases like the murder of Garfield where there is considerable distance between the action and the result, and they have been turning to

lexical semantics for assistance since the very first decompositional theories of word meaning. One of these early theories was the Generative Semantics theory developed by linguists James McCawley and George Lakoff. Lakoff (1965) claims that the sentence 'Floyd broke the glass' is actually derived from an underlying structure like:

(31) Floyd caused (the glass break)¹⁰

McCawley (1968) and Lakoff (1965) proposed that the semantic structures which compose lexical items are themselves words like 'cause', 'have', and 'go'. This would, for example, explain why 'spray' participates in the locative alternation by claiming that it is composed of other words that are the same general type. Generative Semantics might also claim that 'open' participates in the causative alternation because it is composed of the words 'cause' and 'open'. Most controversially, they even claimed that, contrary to appearances, 'kill' does participate in the causative because it is really composed of the words 'cause' and 'die'.¹¹

One of Fodor's earliest arguments against lexical decomposition (Fodor, 1970) targets the Generative Semantics approach of McCawley and Lakoff. Fodor presents examples where a verb like 'kill' is replaced with the corresponding words that it supposedly contains (lexicalizes), 'cause' and 'die', showing that there are differences in truth-conditions (or perhaps acceptability conditions). Therefore, the two cannot be identical. The best examples of this are the following:

- (32) a. Floyd caused the glass to melt on Sunday by heating it on Saturday.
 b. *Floyd melted the glass on Sunday by heating it on Saturday.
- (33) a. Jim caused Bill to die on Friday by poisoning him on Wednesday.
 b. *Jim killed Bill on Friday by poisoning him on Wednesday.

If 'kill' is derived from the words 'cause' and 'die', then we should expect them not to have differences in the conditions that would make them true (or acceptable). Yet in (32) and (33), they do. Therefore, the thesis of Lakoff and McCawley is false. Fodor (1970, p. 6) explains: 'The point is, roughly, that one can cause an event by doing something at a time which is distinct from the time of the event. But if you melt something, then you melt it when it melts'.

It is true that some decompositional theories of lexical semantics assume that the elements in semantic structures are the same type as the words they attempt to explain. These include the original Generative Semantics theories as well as the

¹⁰ It should be pointed out that McCawley and Lakoff were not directly drawing on verbal alternations as data but rather on ambiguities. Positing such an underlying structure to words helps resolve ambiguity in sentences like: 'Floyd melted the glass and it surprised me', which can be interpreted as either the melting itself being surprising, or Floyd doing it being surprising.

¹¹ There is actually some independent support for this from etymology: Pietroski (2003) reports that the Old English words for 'kill' and 'die' were 'cwell' and 'cwell', respectively.

natural semantic metalanguage approach of Anna Wierzbicka (1996). Fodor's criticisms are a serious challenge to these theories. However, the Neoclassical Theory of concepts presented here insists that elements like CAUSE are not of the same general type as the words they attempt to explain.¹² Speakers may indeed pair words like 'kill' with different truth-conditions than phrases like 'cause to die', and they are also likely to have different grammatical properties. This is because 'kill' is not composed of the word 'cause', but the hypothetical structure CAUSE.

To learn about CAUSE, we try to describe some common denominator that verbs like 'clatter', 'rattle', 'break', 'open', 'boil', and 'march' all have in common, and verbs like 'rumble', 'holler', 'cry', and 'walk' all lack. One proposal is that in the causative verbs, some external entity is bringing a change of state in another entity by 'immediate control' that we call CAUSE (Jackendoff, 1983; Levin *et al.*, 1997). Another is that there is some sort of change of state involved. What is the difference between this change of state and the change of state in verbs participating in the locative alternation? We might use the entity INCH, short for 'inchoative', (already posited to explain the conative alternation) to represent 'becoming' a state as opposed to simply going to it.

If we are using the causative alternation to discover more about CAUSE, it turns out that the verb 'cause' is not even a causative verb! At first blush this might sound like Marx's claim that he was not a Marxist. But perhaps it is worth asking why we are interested in causation. Is it to learn more about how we use the verb 'cause', or is it to learn about words we sometimes apply it towards, like 'kill'? If the latter, then we need not even worry about the English word 'cause' at all! This alone is progress in the philosophical debate. Equipped with this progress, we are able to ask (like Dennett on free will): what kind of causation are we interested in? The kind implicit in our inferences using 'kill', 'boil', and 'break', or some metaphysical variety perhaps implicit in our best scientific theories? It may be that there will not be a unified theory of CAUSE and causation, and that the former is not even a real property of the world at all, but a mere projection of our conceptual structure.

Ultimately, how much we can learn from neoclassical conceptual analysis is an open question. It may turn out that not much can be illuminated from non-syntactic linguistic constructions, but the above prospects at least provide some motivation for pursuing the possibility. If successful, the Neoclassical Theory may even help to unify some of the recent work done in analytic epistemology and meta-ethics under a guiding research program.

*Department of Philosophy
University of Pittsburgh, Johnstown*

¹² The most obvious difference is that the elements I am proposing do not have any phonological information paired with them. This is not a strange thing to see in linguistic theory; there are plenty of syntactic structures (like traces and movement) that are phonologically null.

References

- Akmajian, A., Demers, R. A., Farmer, A. K. and Harnish, R. M. 2001: *Linguistics: An Introduction to Language and Communication*. Cambridge, MA: MIT Press.
- Barsalou, L.W. 2008: Grounded cognition. *Annual Review of Psychology*, 59, 617–45.
- Bloom, P. 2004: *Descartes' Baby*. New York: Basic Books.
- Carey, S. 1985: *Conceptual Change in Childhood*. Cambridge, MA: MIT Press.
- Chomsky, N. 2000: *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Cohen, S. 2001: Contextualism defended: comments on Richard Feldman's 'skeptical problems, contextualist solutions'. *Philosophical Studies*, 103, 87–98.
- Cummins, R. 1996: Systematicity. *The Journal of Philosophy*, 12, 591–614.
- Dretske, F. 1983: *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Fillmore, C. 1970: The grammar of hitting and breaking. In R. Jacobs, and P. Rosenbaum (eds), *Readings in English Transformational Grammar*. Washington DC: Georgetown University School of Language, 120–33.
- Fodor, J. 1970: Three reasons for not deriving kill from cause to die. *Linguistic Inquiry*, 1, 429–38.
- Fodor, J. 1975: *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Fodor, J. 1987: *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. 1998: *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J. and Lepore, E. 2002: *The Compositionality Papers*. Oxford: Oxford University Press.
- Frege, G. 1884: *The Foundations of Arithmetic*. Oxford: Blackwell.
- Geeraerts, D. 2010: Theories of lexical semantics. In D. Geeraerts, and H. Cuyckens (eds), *The Oxford Handbook of Cognitive Linguistics*. Oxford: Oxford University Press.
- Hale, K. and Keyser, S.J. 1987: A view from the middle. *Lexicon Project Working Papers* 10. Cambridge, MA: Center for Cognitive Science at MIT.
- Hampton, J. 1979: Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 18, 441–61.
- Hampton, J. 1998: Similarity-based categorization and fuzzyness of natural categories. *Cognition*, 65, 137–65.
- Higginbotham, J. 1989: Elucidations of meaning. *Linguistics and Philosophy*, 12, 465–517.
- Horwich, P. 1992: Chomsky versus Quine on the analytic/synthetic distinction. *Proceedings of the Aristotelian Society*, 92, 95–108.
- Jackendoff, R. 1983: *Semantics and Cognition*. Cambridge, MA: MIT Press.
- Jackendoff, R. 1990: *Semantic Structures*. Cambridge, MA: MIT Press.
- Jackendoff, R. 2002: *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.

- Johnson, K. 2004: On the systematicity of language and thought. *Journal of Philosophy*, 101, 111–39.
- Johnson, K. 2008: An overview of lexical semantics. *Philosophy Compass*, 3, 119–34.
- Kant, I. 1781: *The Critique of Pure Reason*. Cambridge: Cambridge University Press.
- Keil, F. 1989. *Concepts, Kinds, and Conceptual Development*. Cambridge, MA: MIT Press.
- Komatsu, L. 1992: Recent views of conceptual structure. *Psychological Bulletin*, 112, 500–26.
- Lakoff, G. 1965: On the nature of syntactic irregularity. PhD thesis, Harvard University.
- Lakoff, G. 1987: *Women, Fire, and Dangerous Things: What Categories Reveal About The Mind*. Chicago, IL: University of Chicago Press.
- Laurence, S. and Margolis, E. 1999: Concepts and cognitive science. In S. Laurence, and E. Margolis (eds), *Concepts: Core Readings*. Cambridge, MA: MIT Press, 3–82.
- Levin, B. 1993: *English Verb Classes and Alternations*. Chicago, IL: University of Chicago Press.
- Levin, B. and Rappaport-Hovav, M. 2005: *Argument Realization*. Cambridge: Cambridge University Press.
- Levin, B., Song, G. and Atkins, B. 1997: Making sense of corpus data: a case study of verbs of sound. *International Journal of Corpus Linguistics*, 2, 23–64.
- Locke, J. 1690: *An Essay Concerning Human Understanding*. Harmondsworth: Penguin Books.
- Lyons, J. 2005: Representational analyticity. *Mind & Language*, 20, 392–422.
- Machery, E. 2009: *Doing Without Concepts*. Oxford: Oxford University Press.
- McCawley, J. 1968: The role of semantics in a grammar. In E. Bach, and R. Harms (eds), *Universals in Linguistic Theory*. New York: Holt, Rinehart and Winston, 124–69.
- McCloskey, M. and Glucksberg, S. 1978: Natural categories: well-defined or fuzzy sets? *Memory and Cognition*, 6, 462–72.
- Millikan, R. 1984: *Language, Thought, and other Biological Categories*. Cambridge, MA: MIT Press/Bradford Books.
- Mohanan, T. and Mohanan, K.P. 1999: On representations in grammatical semantics. In T. Mohanan, and L. Wee (eds), *Grammatical Semantics: Evidence for Structure in Meaning*. Stanford: CSLI Publications, 23–76.
- Murphy, G. 2002: *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Pagin, P. and Westerståhl, D. 2010: Compositionality: Arguments and problems. *Philosophy Compass*, 5, 265–82.
- Pietroski, P. 2005: *Events and Semantic Architecture*. Oxford: Oxford University Press.
- Pietroski, P. 2003: Small verbs, complex events: analyticity without synonymy. In L.M. Antony, and N. Hornstein (eds), *Chomsky and His Critics*. Malden, MA: Blackwell, 179–214.
- Pinker, S. 1989: *Learnability and Cognition: The Acquisition of Argument Structure*. Cambridge, MA: MIT Press.

- Pinker, S. 2007: *The Stuff of Thought: Language as a Window into Human Nature*. New York: Viking Press.
- Prinz, J. 2002: *Furnishing the Mind: Concepts and their Perceptual Basis*. Cambridge, MA: MIT Press.
- Putnam, H. 1962: It ain't necessarily so. *Southern Journal of Philosophy*, 59, 658–71.
- Putnam, H. 1975: The meaning of 'meaning'. In *Mind, Language and Reality: Philosophical Papers Volume 2*. Cambridge: Cambridge University Press, 215–71.
- Quine, W. V. O. 1951: Two dogmas of empiricism. *Philosophical Review*, 60, 20–43.
- Rey, G. 1993: The unavailability of what we mean: a reply to Quine, Fodor, and Lepore. *Grazer Philosophica*, special issue, 61–101.
- Rey, G. 2008: The analytic/synthetic distinction. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2010 Edition). Available at: <http://plato.stanford.edu/archives/win2010/entries/analytic-synthetic/>
- Rey, G. 2010: Concepts versus conceptions (again). *Behavioral and Brain Sciences*, 33, 221–2.
- Rosch, E. 1975: Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192–233.
- Russell, G. 2007: The analytic/synthetic distinction. *Philosophy Compass*, 5, 712–29.
- Searle, J. 1980: The background of meaning. In J. Searle, F. Kiefer, and M. Bierwisch (eds), *Speech Act Theory and Pragmatics*. Dordrecht: D. Reidel.
- Smith, E. and Medin, D. 1981: *Categories and Concepts*. Cambridge, MA: Harvard University Press.
- Stanley, J. 2004: On the linguistic basis for contextualism. *Philosophical Studies*, 119, 119–46.
- Szabó, Z. 2007: Compositionality. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2010 Edition). Available at: <http://plato.stanford.edu/entries/compositionality/>
- Wierzbicka, A. 1996: *Semantics: Primes and Universals*. Oxford: Oxford University Press.
- Wittgenstein, L. 1953: *Philosophical Investigations*. Oxford: Wiley-Blackwell.