

Georges Rey, Philosophy, Univ of Maryland, College Park MD 20742
until Jan 2004: Philosophy, Stanford University, Stanford CA
georey@earthlink.net; (415)-970-1063; (202)-331-9267

Fri 17 Dec 04 Final Version

limit 30pp, @250wds/pg = 7.5K; here: 8.8K words w/refs; 7.5 w/out

**Philosophical Analysis, Cognitive Psychology,
and The Importance of Empty Concepts¹**

note: square bracketed expressions refer to the concept expressed by the expression within the brackets: e.g. `[ghost]` refers to the concept expressed by the term `ghost`.

Abstract

The "Classical" philosophical view of concepts as having necessary and sufficient defining conditions has fallen on hard times, supposedly both in philosophy and psychology. However, I will argue that both its character and its motivation have been misunderstood: it is not committed to definitions being what people ordinarily use to categorize, or to being otherwise superficially available to their minds. Nor is it committed to excluding prototypes or theoretic roles as possible definitions. What it *is* committed to are claims about the conceptual connections that underlie an agent's competence to make *modal* judgments about what *would* satisfy the concept in actual and counterfactual circumstances, and related facts about what an agent finds *intelligible*.

Why believe there are such conceptual connections? Most discussions focus on ascriptions of concepts of *existent* phenomena -- plants, animals, artifacts-- which permit an ordinary "existential" construal of those ascriptions, and a consequent misleading reliance on actual phenomena in identifying the concept. "Empty" concepts of *non-existent* things --e.g, ghosts, devils-- force us to a *purely intentional* construal, and thereby, I argue, to a consideration of concepts proper, and the need then for classical analyses.

On the other hand, many philosophers and psychologists have underestimated the challenges that have been raised against analyses, by Quine and most recently by Fodor. I review those challenges, but conclude that, as serious as some (though not all) of them are, Quine and Fodor fail to provide adequate alternative explanations of the data that analyses would explain. Analyses seem to be needed despite not being as available as many philosophers and psychologists have hoped. They may be as hidden as the principles of a Chomkian grammar.

1. Delivered as opening address for conference on cognitive science at the University of Quebec at Montreal, June 03. A few short passages have been lifted from my (1993), (1998) and (forthcoming-b).

Introduction

I was asked to present some of the present philosophical perspective on concepts. With no pretense to do justice to the wide diversity of recent work, I thought I might say something on behalf of the much maligned "Classical View" associated with traditional philosophy, showing how it's really no worse off, and actually maybe a little better off than its rivals. In §1, I'll discuss the view's present miseries. In §2 I'll discuss some surprisingly confusing terminological issues that are crucial to understanding any talk of concepts, and especially to assessing the merits of currently fashionable "externalist" theories of conceptual content, which I'll discuss in §3. Once those issues are sorted out, empty concepts -or concepts that apply to nothing, such as [ghost], [phlogiston] and [Euclidean line]- can be seen to present a serious problem for externalism. They, as well as the standard semantic intuitions regarding conceptual connections that I'll discuss in §4, provide evidence for the internal conceptual roles sought by the Classical View that, as yet, has not been accommodated by any of its rivals.

1. Misadventures of the Classical View

Philosophical analysis has fallen on hard times. The "classical view" of concepts on which it was based, according to which concepts have necessary and sufficient "defining" conditions, is widely thought to have been refuted by the simple failure to provide any convincing examples. Worse, through the years it has often been burdened by empiricists, Positivists and a surprising number of psychologists with a program of "reduction" of all concepts to sensori-motor ones. This was linked to the strong "verificationist theory of meaning" that

Quine (1953b) and Putnam (1962/75) rightly pointed out was seriously at odds with good scientific practice. Wittgenstein (1953:\$66ff) further muddied the waters by making much of the fact that people don't seem to have available to them analyses of most of the ordinary concepts they are competent to use, but seem to know merely about ordinary usage and "family resemblances" among exemplars. This spawned several decades of research by psychologists on "prototype" and "exemplar" theories of concepts that were widely advertised as alternatives to the Classical View. And then things took an even grimmer turn when Putnam (1965, 1975) and Kripke (1972/80) pointed to how the analyses of "kind" terms and concepts, like [water] or [polio], seemed not to be of the *a priori* "conceptual" kind philosophers might provide, but were to be supplied by scientists of the respective domain.² Most recently, Jerry Fodor (1998, 2003) has argued that belief in any kind of analytic, conceptual connections sought in the Classical View is where cognitive science --and indeed the whole of the 20th Century-- "went wrong." He joins others in advocating an "externalist" theory, whereby the content of a concept is constituted not by any role internal to a thinker, but by the worldly phenomenon to which it bears some real causal relation (§3 below).

Well, despite all this misery for the Classical View, the theme I want to pursue here is that much of it is unwarranted, arising from confusions about what the proper role of analyses were supposed to be, and from burdening that role with conceptions of the project of analysis that were not only needless, but independently bad ideas. I think there are still good reasons to believe in analyses, even if

2. See Rey (1983/99, 1986) for discussion of all these problems.

philosophy may have to engage cognitive scientists –and cognitive scientists, philosophers– to discover them.

What *is* crucial to the Classical View are the conceptual connections underlying an agent's competence to make *modal* judgments with a concept: what she thinks *would* satisfy a concept in various actual and *counterfactual* circumstances, and related facts about what she finds *intelligible*. The existence of these connections is not in the least undermined by the failure of verificationism, the superficial unavailability of analyses, or people's exploitation of prototypes. Whether or not someone reasons prototypically, if on *reflection* she didn't find the possibility of, say, a female doctor in a black coat even *intelligible*, but as unintelligible as "married bachelor" or "round square," that would be a reason to think she didn't really have the concept [doctor] –that 'doctor' for her meant maybe [typical American doctor]. This role for concepts seems to me to be related to one of the main advantages of conceptual ascription in the first place: it identifies what agents can *understand*, and how their attitudes might in principle be modifiable by evidence and argument, abstracting *semantic* issues from the *epistemic* ones with which they are too often conflated. More colloquially: it factors out the "merely verbal" issues from the genuinely substantive ones.

An Externalist like Fodor, however, might argue that questions of modality and intelligibility are questions a person settles not by consulting meaning-constituting analyses, but simply by learning a good theory of the actual phenomenon in the world to which a concept refers. Although, for reasons we will explore in §§2–3, this option can seem plausible in the case of concepts that do genuinely refer, it is obviously problematic for *empty* concepts, or concepts that don't

refer to any real, or even genuinely possible phenomenon in the world, for example, [ghost], [demon], [bogeyman], [phlogiston] and the like. And, as we shall see in §3, empty concepts are by no means confined to fairy tales and bad science: they may be a crucial part of our normal, successful interaction with the world (Brentano (1874/1973) famously argued that they were a distinctive mark of the mental). I'll argue in §4 that a psychology of such concepts, and of the rich mental and behavioral processes to which they give rise, requires something at least very like the analyses that the Classical View sought.

On the other hand, there is no question that many philosophers and psychologists have underestimated the difficulty of producing these analyses. For all my belief in analyses, I am impressed by the sceptical challenges that have been raised against them. Consequently, in my remarks here I want to engage in a rather delicate dialectical balancing act: on the one hand, urging the need for philosophical analysis in any cognitive psychology, particularly one treating empty concepts; but, on the other, drawing attention to the formidable problems of providing any. This, at any rate, ought to provide some idea of how the topic presently stands in philosophy.

2. Terminological Issues

Before proceeding, I want to address some surprisingly confusing terminological issues regarding concepts that deserve far more attention than they've received.

2.1 Existential vs. (Purely) Intentional Usage

Talk of empty concepts is talk of thoughts, ideas, representations that are "about nothing." But there are two ways that a representation could be about nothing: it could be *meaningless*, as in the case of nonsense expressions, like 'brillig'. Or it could be perfectly meaningful, but, as in the case of 'Zeus', there be no *real thing in the world* that it represents. These two ways of being "about nothing" give rise to a crucial ambiguity in ways of talking about "what representations represent." On the surface 'represent' would appear to be simply a two-place relation, as in:

(1) The word 'cats' represents cats.

But this can't be quite right, since

(2) The word 'elf' represents elves

would then be false, for lack of elves: you can't bear a real relation to something that doesn't exist. But there's surely a reading of (2) that makes it true, since, again, 'elf' is not *meaningless*. It's merely "empty" (which I'll confine to meaningful expressions).

So what does an "empty" term like 'elf' represent? Well, it's an interesting fact that an almost universal response is that of Quine's (1953a) fictitious philosopher, "McX": it represents "an idea in your head." Now, on the face of it this would seem absurd, since (a) whatever else might be in your head, there are certainly no elves there; and, in any case, (b) if elves are actual ideas in your head, then *elves would turn out to exist after all!*

I won't go through all the replies that have been made to this puzzle, from Plato past Quine (see esp Cartwright 1960/87), but rather simply note that the word 'represent' (and, for that matter, virtually any intentional idiom) seems to suffer from a systematic ambiguity

between along the following lines:

(REP) (i) If we are talking about a representation, *x*, of some *real thing y*, then *x* represents that real thing *y* -thus 'Nixon' represents the actual man Nixon.

(ii) When there isn't, as in the case of 'Zeus', then we rely on talk about the content of the expression 'y' (which I will abbreviate by placing brackets around an expression, e.g. [Zeus]).

I'll call the first usage the "existential," the second the "(purely) intentional" usage of 'represent' (and other intentional idioms).³ I've expressed the second, intentional use with deliberate vagueness. It would be tempting to say "so a purely intentional use of 'representation of y', for lack of any *y*, is really about an intentional content." But this wouldn't be correct, since someone thinking about Zeus and his philandering ways isn't thinking about the philandering ways of an intentional content. Speaking more carefully, we should say something like: when *x* represents a *y* that doesn't exist, a person is standing in the thinking relation to [*y*]; but this doesn't entail she is thinking about [*y*]. Even in a purely intentional usage, "thinking about *y*" is one thing; "thinking about [*y*]" quite another (see Rey (in prep) for further discussion).

3. The distinction is, of course, close to the much-discussed distinctions between "transparent"/"de re" and "opaque"/"de dicto" readings of propositional attitudes and/or their ascriptions (see, e.g. Kaplan 1969). I don't want to assimilate my distinction immediately to those, both because they are the objects of enough controversy on their own, and because a usual strategy - e.g. Kaplan's--for understanding them won't work for 'represent': e.g., a transparent reading of "John thinks of Sam Clemens that he's funny" may well involve John being related to a representation, "Mark Twain is funny," that involves a representation, "Mark" that in fact represents Sam. But this understanding obviously can't be available for the term 'represent' itself.

Interestingly enough, an analogous ambiguity can arise for the preposition `of' in `concept of'. For example, so long as (it's believed that) there's a referent to a concept, it's easy to rely on it to identify the content of the representation. Thus, we cheerfully say of the Ancient Greeks, "They thought the stars were holes in the heavenly canopy," not flinching at the extraordinary fact that this means the Greeks didn't even think stars were material objects! By contrast, I suspect we would flinch at "The Greeks thought that Mercury was an angel and that Pluto ruled hell": "Angel"? "Hell"? Did the Greeks really have *these* concepts?! Arguably not. But their beliefs about Mercury and Hades were probably closer to our (conditional) beliefs about angels and Hell ("If there were angels...") than their beliefs about the holes in the canopy were to our beliefs about stars. Nevertheless, the Greeks presumably had no concept [hell] --and so perhaps they didn't really have a concept [star] either, and we've merely lazily relied on existential usage in saying that they did.

Some might allow that the Greeks had "a very different concept of star" than we have. This would appear to suggest, incoherently, that the Greeks both did and did not share with us a concept with the content [star]. But arguably such talk is an attempt to combine an existential and intentional ascription to the Greeks: they represented *the stars*, but they did this other than by using the concept [star]. Of course, this invites the question of why we should think that whatever concept they did have -[holes in the heavenly canopy]?- really did represent the stars. `Concept of x', especially used existentially, when there is a real x on which we can rely, tends to obscure this crucial question.

2.2 Concepts as Between Representations and Referents

Another way that that question can get obscured is by claiming that, e.g., the Greeks simply deployed a *different representation* of the stars, much as different people might have different photographs of them. This accords with much psychological usage: if we're investigating a child's "concept of a bird," we often uncritically take for granted that the child and we are both thinking about these real things, birds, and we then concentrate on the particular way in which these birds get represented, e.g. by "prototypes" by "theories," or by definitions, where these are understood to be items entokened in the brain. I think this in part is what leads psychologists in particular to use the word 'concept' for merely *internal mental representations*.

However, mental representations would no more be identical to concepts than are the words in a natural language. Words in a language are usually individuated *syntactically*, in such a way that different syntactic types, e.g. the words 'city', 'metropolis', 'ville', 'Stadt' all arguably express the *same concept*, as might pictures of New York or Paris. In the case of mental processes, one person might express [number] thoughts by decimal numerals, another by binary ones; and, on some views, others might represent them by images, spatial layouts or by parts of the body. Moreover, *the same representation* may at different times or by different people express different concepts: an image of a body part might represent not a number but that very body part!⁴ Consequently, there must be something about the role of the representation, either in relation to its referent or in relation to other representations, that determines

4. Talk of mental imagistic representation does seem to me, however, fraught with confusions. See Pylyshyn (2003) and Rey (1981, in prep) for discussion.

which concept it expresses. This is the notorious third "semantic" realm, over which theorists contend. I will argue that philosophical analyses are required not for an account of *the computational character* of representations, but for an account of *the conditions of conceptual competence*, i.,e. for someone's having a representation with a *specific conceptual content*.

3. The Inadequacies of Externalism

So what are those conditions? Well, there's the "Classical View" according to which a representation expresses a certain content that can be identified by a set of necessary and sufficient defining conditions. That's the view of "philosophical analysis" that's fallen on hard times and to which I will return after considering a recently influential alternative to it, viz. the "Externalist" approach that has arisen in the last thirty years in response to the aforementioned work of Putnam and Kripke.

In a way, Externalism might be said to make a principled virtue of what I have been identifying as the above reliance on merely representations and referents, dividing all the work of concepts between the two. In particular, there is what I will call the *Strong Externalism* advocated by a great diversity of recent philosophers: e.g., Stampe, Dretske, Stalnaker, Fodor, Millikan --and, before them, Skinner and Quine.⁵ To a first approximation (the details and many

5. See Rey (1997:ch 9) for a review of these proposals. It is worth noting that Strong Externalism is also taken for granted even by theorists who *don't* believe there are real phenomena in the world that representations represent. Thus, Chomsky (2000) emphasizes how phonetic phenomena don't really exist in the acoustic stream, but, amazingly enough, this leads him to suppose that representations don't have any intentional content at all! I suspect this extreme reaction is due in part to his presuming, with many of his opponents, that any intentionalism must be Strongly Externalist (see Chomsky (2000:156-60) and Rey (2003) for discussion).

differences between their views won't matter here), x has the concept y iff some state of x stands in the right causal relation to some actual y: e.g. would causally co-vary with y, was caused by y under certain (ideal, normal, evolutionarily significant) conditions. Thus, for example, someone has the concept [horse] iff she could, under certain conditions, tell the horses from the non-horses; or has a state that was evolutionarily selected for responding to horses.⁶ The important feature of Strong Externalism is the claim that the content of a representation is constituted by the real phenomenon to which it is, one way or another, causally related. Different concepts of the same referent are simply different representations, often with different logical structure, that are causally linked to the same phenomenon.

I hasten to distinguish this Strong from *Weak Externalism*, according to which the content of *some* representations depends *in some way or another* upon *some or other* interesting relations the system of which they are a part bears to an external environment. This view allows that ascription of content to a system cannot in general be done in complete abstraction from the world in which the system happens to be embedded, but leaves wide open the specific ways in which that embedding may figure in any particular ascription. Perhaps a direct causal route between representation and referent is required for perceptual demonstratives; perhaps some sort of co-variation for certain predicates; and who knows what for fictional and other empty terms. I won't dispute this weak view here. It seems to me supported

6. Or, on Fodor's (1991) ingenious "asymmetric dependency" (or "locking") proposal, she would token 'horse' in the presence of non-horses only if it was a law that she would token it in the presence of horses, but not *vice versa*. The complexities of this account won't be an issue here.

by the usual examples --("twin" cases as in Putnam (1975) and Burge (1979)-⁷ and, indeed, is pretty much *all* that is supported by them: Strong Externalisms are (to my mind) tenuous *theories of meaning* intended to *explain* them.

One kind of problem faced by Strong Externalism is how to specify the right causal relation without presupposing the very intentionality one is trying to explain.⁸ Although I suspect that is a likely insuperable problem, I don't want to pursue it here. Rather I want to consider the empty concepts that pose an obvious problem for any Strong Externalism, even one specifying the right causal relations. Whatever those relations might be, they can't be had to things that don't exist.

Fodor (1987, 1991, 1998) has been the most explicit about how to deal with such cases, making a familiar, but what seems to me excessively easy move. True, enough, there are no elves; but, he claims, there *might have been* --there's the (uninstantiated) *property of elfhood*. And, for Fodor, it's enough that there's a law that 'elf' *would* be caused by elfhood *counterfactually* for the right meaning-constituting causal relation to be secured. What Fodor is exploiting here is a certain profligacy about properties that isn't usually tolerated for mere material objects: for Fodor, properties can exist *uninstantiated*, and laws involving counterfactuals, about what *would*

7. These are the examples of beings that are physical duplicates, but whose thought contents seem to depend upon which environment they inhabit. The simplest case is provided by proper names: imagine a duplicate of Bush inhabiting a planet that enjoyed precisely the same kind of physical history as the earth; the proper name thoughts of his twin there would clearly not be representing the very same people or places in *his* life as Bush's ordinarily do in his (the twin's "Iraq" thoughts would refer to an entirely different place from Bush's), despite their internal physical type identity.

8. See the introduction and essays in Loewer and Rey (1991), as well as Greenberg (2000) for various criticisms along these lines.

have happened if the world had been different than it is, are apparently there just for the asking.

Against this view, I would like first to enter a certain methodological protest, or anyway, a plea for honest toil: *except as a last resort, do not introduce into your psychological ontology entities for which there is no independent evidence.* Before we glibly resort to *elfhood* --or to *being phlogiston*, or *being divine*-- and to laws and counterfactuals regarding them, we need to ask whether, independent of our semantic desperation, there is any reason whatsoever to believe in such things? I know of none. Our best theories of the world --biology, chemistry, physics, cosmology-- do not seriously make any room for such things or properties, or laws or counterfactuals relating them.

Nor does anyone seriously expect them to. It's not that natural scientists might not have occasion to seriously postulate *some* uninstantiated properties. There may be many instances of magnitudes specified by differential equations, or isotopes, transuranial elements or other phenomena that never happen to be instantiated in the actual history of the world, but which nonetheless need to be included as values of physical variables. It's just that ghosts, elves and unicorns are surely not among them. Indeed, note that, revealingly, whereas we would turn to *chemists and physicists* to learn about the uninstantiated magnitudes or elements, we would turn not to them, but to psychologists or folklorists historians to learn about, well, of course, not elves or ghosts *themselves*, but about people's *thoughts* about them --reflecting precisely the switch in empty cases from existential to purely intentional usage of 'represent', which, we observed earlier (§2.1), directs us to thinking about concepts proper,

rather than their referents.

In any case, *some* properties are, uncontroversially, *necessarily uninstantiated*, e.g. [round square], [largest prime], [red and not red], [round-square]; and if a property is necessarily uninstantiated then surely there *cannot* be laws with which a representation can causally interact, actually or counterfactually. Fodor (1990, 1998) briefly addresses this problem, and, observing that many examples are logically complex, articulates a famous view that goes back to Russell (1912/59): "There can be no *primitive* concept without a corresponding property for it to [covary with]" (1998:165). Fodor simply expects there is metaphysics enough for all our primitive thoughts, an instantiable property for every primitive predicate. This seems to me a bit rash. Indeed, there seem to me to be lots of cases that would at least appear to present *prima facie* counterexamples.

Take, for starters, the familiar primitive concepts of Euclidean geometry that most high school students grasp, e.g. [point], [line], [plane], [circle]: as many since Plato have pointed out, these concepts may well be *uninstantiable*. At any rate, we could never be sensorily *presented* with any such thing --all perceptible points and lines have some thickness, and so no representation in our head could enter into causal relations with any such thing or property.

Nor are uninstantiable primitives limited to "geometrical" examples. There is a long tradition of philosophers --Spinoza, Hume, Wittgenstein-- who have plausibly argued that nothing could possibly satisfy such primitive "supernatural" concepts as [soul], [elf], [miracle], [magic], [monster], [free will], [destiny], [karma] (see Slote 1975 for discussion). And, moving away from these supernatural examples, many philosophers have raised analogous worries about our

ordinary folk concepts of [color], [sound], [cause], [material object].

In these latter cases, philosophers are often spontaneously joined by psychologists, who are fond of claiming that colors, sounds, phones and phonemes don't really exist: these things are (as they often put it, echoing McX) "merely in the mind." In their classic discussion of English phonology, for example, Chomsky and Halle (1968) write:

there is nothing to suggest that these phonetic representations also describe a physical or acoustic reality in any detail. For example, there is little reason to suppose that the perceived stress contour must represent some physical property of the utterance in a point-by-point fashion... In fact there is no evidence from experimental phonetics to suggest that these contours are actually present as physical properties of utterances in anything like the detail with which they are perceived. --Chomsky and Halle (1968:p25)

And, in a recent, comprehensive text on vision, Stephen Palmer (1999) writes:

Neither objects nor lights are actually "colored" in anything like the way we experience them. Rather, color is a psychological property of our visual experiences when we look at objects and lights, not a physical property of those objects or lights. ... Color is more accurately understood as the result of complex interactions between physical light in the environment and our visual nervous systems. ... There may be light of different wavelengths independent of an observer, but there is no color independent of an observer, because color is a psychological phenomenon that arises only within an observer. --Palmer (1999:97-9)

If these views are correct, then even the content of the primitives of linguistic and visual perception cannot be constituted by their

referents in the real world.⁹

At any rate, examination of many of these concepts reveals that people on reflection are hard put to say what a world would be like in which there were real colors, Euclidean shapes, people were relevantly free to do otherwise, and souls could survive bodily death. Quite apart from the metaphysical issues, *these facts are by themselves psychologically important ones needing explanation. It's hard not to think some of this persistent puzzlement is due to there being something inherently problematic about these concepts. If there are at least two such cases, an Externalist theory will face a difficulty distinguishing them.*¹⁰

Fodor deals with empty cases by being generous with the (possible) phenomena the world provides. Another externalist, Ruth Millikan (1998) takes a different, more niggardly line: she denies there are actually *any genuine thoughts* at all in such cases! She writes that such empty substance concepts

are not substance concepts at all. An ability that is not an ability to do anything is not an ability at all. Empty substance concepts result from failures of the mechanisms designed develop

9. Neither Chomsky and Halle, nor Palmer, of course, address the (peculiarly philosophical) modal question of whether phones or colors *could* exist in some other *possible* world, but the burden is surely on the externalist to show how they could, especially in view of the difficulties of identifying them in the actual one. Fodor (1998) suggests a dispositional strategy, but this risks rendering his Externalism circular (see Rey (in prep) for discussion).

10. In reply to my challenging him (in my 2003) on why we should expect the world to so conveniently conform to our thought as to supply a property for every primitive predicate, Fodor (2004) curiously replies:

What I suppose is that our ways of thinking about [the world] accord pretty well with the way the world is. That's sort of what you'd really expect. Making thoughts that accord with the world is what cognition is *for*. -(p109fn13)

What is curious about this reply is that it goes so much against the grain of his (1987) otherwise quite reasonable scepticism about teleological and/or selectionist accounts of psychological traits. Indeed, as examples of geometrical contents in the visual system show, a system with necessarily empty concepts may get on very well -provided, I suppose, that the errors it makes don't matter.

substance concepts. They are "concepts" only in that their biological purpose was to have been concepts. --Millikan (1998:175-6)

When pressed for clarification, she insists that empty terms are not "real terms," that they are like "an ability that is not an ability to do anything," that it "makes no sense to talk of a mode of presentation that doesn't present anything," and that, indeed, when reference "fails utterly there is no thought at all" (pc). This seems to be a generalization of a view Gareth Evans (1982) sketched regarding *empty singular terms in natural language*, but much more radical in ways that I doubt he himself would have endorsed.¹¹

In the first place note that, per my earlier terminological observations, it makes perfect sense to talk of "a mode of presentation that doesn't present anything": it's just that one needs to distinguish, as I have, existential and intentional usages. More importantly, if apparent "thoughts" with empty terms are not really thoughts at all, then how on earth are we to explain the often rational, *content sensitive* behavior of people (and their subsystems) that seem to have them? People pray, make sacrifices, and engage in often elaborate reasonings about gods, devils, elves, angels, ghosts; the visual system seems to compute over representations of Euclidean figures and colors; and astronomers once intelligently sought evidence for the luminiferous ether and the planet Vulcan. In all such cases, empty thoughts interact in myriad *inferential* ways with non-empty ones (e.g. about churches, misdeeds, light, the sun), ways that surely require the empty ones to possess *some kind* of intentional content.

11. But even Evans' and related views (e.g. Taylor's (2003), Campbell's (2002)) would seem to face serious *prima facie* problems with our apparent ability to refer to stable, illusory "things" such as "particular" rainbows, animated cartoon characters and Kanizsa triangles; see Rey (forthcoming-a and in prep) for discussion.

4. The Need for Internal roles

I don't see how we can hope to do the psychology of such empty thoughts, particularly the necessarily empty ones, without saying something about the role of these thoughts in a person's mind: again, if you want to know about the nature of *ghosts*, you look not to what people have been "getting at" *in the world*, or even in any genuinely *possible* world, but to merely what they *think*. [Ghost], but not [will], is tied to the imagined possibility of a disembodied mind; [will], but not [ghost] is tied to maybe impossible claims about spontaneity and moral responsibility; Euclidean 'point', 'line', 'circle' form a cluster of interdefinable notions subject to certain axioms; 'color' and 'sound' with projections onto the world of certain stable reactions in our visual system.¹²

Such a claim is not without some independent motivation. Internal roles are generally acknowledged to be needed for the concepts of logic and mathematics ([not], [all], [successor]), for personal response dependent concepts such as [funny], [cute], [worthwhile], and perhaps even for the likes of [beautiful] and [good], which also present problems for any purely externalist view. At any rate, it is doubtful there are real phenomena of "ands" and "nots" in the world, or real properties of *being funny* or *cute*; and, even if there were, people with "the same concept" of, e.g., funniness might well disagree irresolvably,

12. In saying that roles here are crucial, I don't mean to be denying a Weak Externalism (§3 above), according to which assigning any content whatsoever to a system may depend generally upon the environment in which it is embedded. Nor do I intend to deny that the content of *some* terms (e.g. 'water') that do manage to refer to something real acquire their content thereby. I am not here in the business of *providing* a theory of content, but merely denying what seem to me rash ones that have been proposed.

and so differ in their causal relations to external phenomena.¹³

4.1 The Quinean Challenge

But now for the other side of the dialectic. As pressing as the need might be for some kind of "role" account of concepts, there are substantial difficulties in providing one. For starters, there's the fact that people seem to be able to believe most anything --Berkeley that trees are ideas, Pythagoreans that everything is numbers, contemporary literary theorists that everything's a text. Moreover, as Quine (1953b) has emphasized, all beliefs seem *revisable*: after all, you can at least fool some of the people some of the time about most anything. So it would appear impossible to tie a concept to any specific role.

However, these considerations are hardly decisive. I have argued elsewhere (Rey 1993, 1998) that Quine's "revisability" argument, like much of the psychology of his time, suffers from an excessive "superficialism," or the bad idea that psychologically real distinctions should be available on our behavioral or introspective surface. People can be mistaken about most anything and revise their beliefs in the light of further experience and argument. *But that doesn't show that there aren't deeper inferential principles governing their thought that might belie these superficial errors.* There is all the difference in the world between the clever Bishop Berkeley, who maintains all objects are ideas, and someone genuinely confused, who really does think without further ado that objects don't persist unperceived. Berkeley, after all, did feel obliged to enlist God for the tree in the lonely forest, exhibiting

13. Besides Jerrold Katz, e.g. his (1972), the philosophers who have recently tried to develop a role theory in most detail are Frank Jackson, in e.g. his (1998), and Christopher Peacocke, in e.g. his (1992, 1998, 2004). See Fodor (2004) and Rey (1996, 1998) for discussion of the latter along lines of the present essay.

his *competence* with [material object], despite his perverse performance. He's simply like someone who deliberately violates some rule of grammar (which is not to say there might not be good reasons for them -or for Berkeley- to do so!).

4.2 The Analytic Data

The analogy with grammar seems to me quite close. As with grammar, we are confronted by regularities that need to be explained. In the kinds of cases that the analytic is supposed to explain, these include, *inter alia*, a non-negligible convergence in people's judgments with regard to particular concepts and/or lexical items, e.g., about synonymies, ambiguities, antinomies, entailments.¹⁴ Ask speakers whether 'bank', or 'pride' are ambiguous; whether killings necessarily involve dyings; whether 'open' and 'closed' are antonyms; whether 'Jim boiled the water' implies 'The water boiled'; whether 'The square-root of a shadow is red' is anomalous. Ask people what things qualify as a bachelor, knowledge, coercion, a voluntary act, perception, water, brisket; elicit their considered judgments about hypothetical cases; ask them what they think *constitutes* being one of these things: I submit that one finds what philosophers have found since Socrates, that there is great deal of convergence about what clearly does and what clearly does not satisfy a concept, and that, although people cannot readily provide adequate rules or definitions to capture these patterns, they soon *realize* they can't, and, despite disagreement, are quick to acknowledge at least *the cogency* of each others' cases and considerations. Whether or not they call them this, or are even justified in

14. See especially Katz (1972:3-6) for a rich survey of the data that seem to me to have been far from sufficiently discussed in the critical literature.

thinking in these ways, people do *seem* to have analytic intuitions, which I call "the analytic data," leaving it open whether it is to be explained by an actual analytic.

Alas, however these analytic intuitions are a lot more slippery and elusive than the lovely examples of bad grammar. The history of philosophy is filled with purported analyses by one generation that have been refuted with counterexamples from the next: there's an entire sub-industry that has been devoted for some time to analyzing "x knows that p," still without success (see, e.g., Morton 1997).

Many of the difficulties in providing analyses, however, *could* be taken as actually an argument *for*, not against, the project of trying to find them, or something like them (it's a little like the *combination* of success and failure in science that argues for realism about its objects: we get enough success to think we're on to something, and enough failure to think that it's not all wishful thinking either). For what is it that leads us to *reject* proposed analyses? Why do we tend to agree about plausible candidates and also about their inadequacies? What do many of us know about 'know' that leads us to fall for 'justified true belief' and then to appreciate Gettier's counter-examples?¹⁵ Or about gods and ghosts to know that they had better *have minds*? Or about "murder" and "person" such that we understand why e.g. abortion and euthanasia are difficult cases? That is to say, we should

15. These are examples against the Classical proposal that [knowledge] is [justified true belief], see Gettier (1963). Nichols et al (2003) present experimental evidence that people don't in fact converge in their reactions to these examples cross-culturally. However, as they acknowledge in response to Frank Jackson, this could be because they are deploying different concepts -just as divergence about grammatical intuitions may be due to idiolectic variation. Indeed, given the role I've emphasized of concepts in guiding modal judgments, this would seem the reasonable conclusion to draw (once one controls also for the hosts of pragmatic influences surely also at work in the interview contexts). The appeal to intuitions is only as data to be explained, even if the data are diverse.

ask here exactly the question that Chomsky asked about syntax: *what explains the patterns and projections in peoples' judgments?*

An interesting case in point is Fodor's (1981) rejection of any analysis of `paint_{tr}` (the transitive verb). He proposes and rejects a number of plausible proposed analyses, ending with:

(P) x paint_{tr} y iff x is an agent and x intentionally covers the surface of y with paint_n and x's primary intention in covering the surface of y with paint is that the surface of y should be covered with paint_n in consequence of x's having so acted upon it. -Fodor (1981:287)

He then raises (p288) still a further counterexample of Michelangelo dipping his *brush* into the paint, and so satisfying (P), but presumably not thereby "painting his brush." At this point, Fodor concludes:

I don't know where we go from here. For all know -for all anybody knows-- `paint_{tr}` is undefinable; for all I know, you can't eliminate it even in terms of such a closely related term as `paint_n'. Or perhaps it *is* definable, but only in a reduction base that includes `dinosaur' and `Chlorodent'. Either way, the present point is that [standard examples] don't work. That's not surprising; when it comes to definitions, the examples always don't work.

--Fodor (1981:288, emphasis his)

Now, first of all, it seems to me odd that Fodor breaks things off at this point -just when things are getting interesting! Several lines of continuation suggest themselves -specifying a still more detailed intention, restricting the scope of an analysis and allowing for pragmatic variation in how `paint' is understood in a context, perhaps understanding analyses as constitutive default rules, to be overridden in certain circumstances.¹⁶ But this is not the place to take up any of

16. This is no *ad hoc* maneuver. A number of philosophers and linguists have pointed to the immense difficulty of distinguishing semantic from pragmatic issues in evaluation of ordinary intuitions of "meaning": see Bach (2001) and Pietroski (2003) for discussion. For default rules, see Thomason (2001).

these suggestions. My points are methodological:

(i) *how and why do we so readily see the plausibility of (P) and the cognency of Fodor's further example? Is it really because (à la Quine) we have such a good empirical theory of the activity of painting? Could we be refuted by discovering that 6 out of 10 painters do paint their brushes? Does anyone care whether Michelangelo was painting his brush -apart from their grasp of the meaning of 'paint_{tr}'? and*

(ii) *why think uncovering the analysis of an item -determining either the primitives or the analysis- ought to be easy? Why not presume it will be at least as hard as uncovering the principles of grammar?*

In any case, *being undefinable* is hardly the natural conclusion to draw from the series of increasingly plausible analyses Fodor considered. We do seem to be converging on something, but, unsurprisingly, are maybe missing some subtle resources.¹⁷

17. Consider, for example, how analyses of material object concepts might be fruitfully informed by the geometric constructions of vision theory (as in Marr 1982, Biederman 1987, Palmer 1999), rather than traditional sense data. And consider the considerable resources of Ramsey sentences in framing definitions in terms of (some selection of) the claims of a theory introducing and relating many terms at once (see Lewis 1972). The opera isn't over until a good many people have sung a lot more arias.

Fodor worries in his discussion of "paint_{tr}" about how "it is vastly unlikely that children have access to [the concept of *primary intention of an act*" and that "the amount of subsidiary apparatus you need to define [the term] is getting sort of hairy" (1981:287-8). But it's hard to see why these should be worries for someone, especially like Fodor (in the same article!), who is prepared to credit neonates with a fairly hairy innate conceptual endowment, especially as regards language.

4.3 Rival Explanations

One way to appreciate the role of some kind of analytic is to consider the alternatives.

(a) Quine's

One of the earliest and most influential alternatives was Quine's own, according to which what was analytic was simply what was centrally or tenaciously believed by someone, capturing the idea that "no amount of evidence could refute the claim that bachelors are unmarried." I won't rehearse all the inadequacies of this account.¹⁸ Suffice it to say that many unrevisable beliefs (e.g. The world has existed for more than ten minutes, There are many people in the world) aren't the least analytic, and many analytic claims (e.g. bachelors aren't married, kissing involves touching, a podiatrist is concerned with feet) aren't the least "central": indeed, the only reason most of us would resist giving up "Bachelors are unmarried" is that *that's simply what the word means* --if someone wants to change the meaning to include women, fine! What's especially striking and interesting about purported analytic claims is not that they are *unrevisable*, but that, as *literally understood* their denials seem *unintelligible*.

18. See my (1992), Fodor (1998).

(b) Fodor's

Fodor (1998) recognizes that Quine's "centrality" explanation won't quite suffice to explain the analytic intuitions,¹⁹ and so follows up a suggestion originally advanced by Hilary Putnam (1962/75): "analytic truths" are just examples of "one-criterion" concepts, or concepts like, e.g. [bachelor], [widow], [ophthalmologist], where there purportedly happens to be only one "way to tell" whether they apply. He acknowledges (p82) that, so stated, this latter account won't suffice either, since the notion of "criteria" seems no better off than "analytic." In particular, if there were just one way to tell what's what, there would seem, trivially, to be indefinite numbers of different ways -for example, just ask someone who knows the one way; or ask someone who knows someone who knows; or... etc.- and we would seem to have no better way to single out which way is "criterial" than we have to say which way is "analytic."

But here Fodor offers a promising move, redolent of his general "asymmetric dependency" theory of content (see fn 6 above): among the ways of telling what's what, some do and some don't asymmetrically depend upon the others; the one(s?) that *don't* are the ones that give rise to analytic intuitions. Thus, telling that someone is a bachelor by checking out his gender and marriage status doesn't depend upon telling by asking his friends, but telling by asking his friends does depend upon telling by his gender and marriage status; and so we have an explanation of why "bachelors are unmarried males" seems analytic

19. At any rate, he (1998) notes it doesn't work "for all the cases," suggesting it works for some. I actually don't find it works for many, except as a matter of pure coincidence, e.g., cases where the term involved just happens also to be central to our thought, as in "Material objects are extended."

-without it's actually being so.²⁰

Such asymmetric dependence will "explain away" the analytic, however, only if the analytic isn't needed in an explanation of such asymmetric dependence. But it's hard to see how it wouldn't be. For notice that not just any such asymmetric dependencies give rise to analytic intuitions: the only way I know to tell whether something's an acid is to see whether it turns litmus paper red; other ways (asking my chemist friends) asymmetrically depend upon that way. But "acids turn litmus paper red" is surely not *analytic*, even for me, since I bet there are, or could be, better ways to test acidity; and that's because 'acid' marks a relatively "deep" sort of thing, whose nature is not captured by what tests happen to be at hand. By contrast, 'bachelor' marks a superficial kind, whose nature is pretty much exhausted by the linguistics of the matter: unlike the case of litmus paper and acidity, *the reason that gender and marriage status are the best way to tell whether someone's a bachelor is that that's just what 'bachelor' means!* Indeed, should a chemist propose revising the test for acids in the light of better theory -perhaps reversing the dependency of certain tests- this would not *per se* constitute a change in the meaning. Should, however, a feminist propose, in the light of better politics, revising 'bachelor' to include women, this obviously would.

In any case, it should be clear that appeals to "one criterion" concepts won't answer the explanatory challenge one might raise for the less obvious cases of the connections between, say, 'know' and 'justified', 'freedom' and 'spontaneity', since in each these cases,

20. Actually, one might wonder why Fodor's proposal is not a reductive explication rather than the eliminative "explaining away" of the analytic that he intends; cf. Horwich (1998), who claims meanings just are the uses of a term that explain all other uses. One reason might be the very reason I will now suggest against it's even being a successful eliminativist alternative.

although we seem to be on to some connections, it's notoriously difficult to specify even *one* genuinely adequate "way to tell."²¹

VII. Conclusion

To summarize: reflection on our capacity to entertain primitive concepts with (necessarily) no referents, forces us to recognize that, *pace* Strong Externalists like Fodor and Millikan (but *modulo* Weak Externalism), some conceptual contents are individuated by some aspect of their role in our thought. Reflection on the variability of the roles of a concept, across people and within a single person across time, forces us to recognize that these roles are not necessarily available in a person's superficial behavior or thought, but must be a part of a background competence. A traditional task of "philosophical analysis" could be regarded as aimed at uncovering those background roles. It is this task that seems to me of a piece with the task of an adequate psychology of concepts, which may, indeed, be as difficult as isolating the principles of a Chomskyan grammar.

21. There's also the problem of how to apply the proposal to many of the obviously empty cases, such as [angel] or [demon]: how, after all, does one actually *tell* whether something's an angel or the devil, as opposed, say, to some misbegotten biology? Or whether someone has an immortal soul? Beats me. But I do presume that it's analytic that angels are some sort of supernatural mental agents, and that souls are the repository of our personal identity. Again, we look not to (how to determine) any angelic properties *in the world*, but merely to how people *think* about such other-worldly matters.

References

Bach, K. (2001), "Speaking Loosely: Sentence Non-Linearity," *Midwest Studies in Philosophy*, vol. 25: *Figurative Language*, P. French and H. Wettstein (eds.), Oxford: Blackwell, pp249-63

Biederman, I. (1987), "Recognition by Components: A Theory of Human Image Understanding," *Psychological Review*, 94(2):115-7

Bonjour, L. (1998), *In Defense of Pure Reason*, Cambridge University Press

Brentano, F. (1874/1973), *Psychology from an Empirical Standpoint*, trans by A. Rancurello, D. Terrell, and L. McAlister; London: Routledge and Kegan Paul

Burge, T. (1979), "Individualism and the Mental," *Midwest Studies in Philosophy*, pp73-121

Campbell, J. (2002), *Reference and Consciousness*, Oxford Univ Press

Cartwright, R. (1960/87), "Negative Existentials," in his *Philosophical Essays*, Cambridge: MIT Press

Chomsky, N. and Halle, M. (1968), *The Sound Pattern of English*, Cambridge (MA): MIT Press

Fodor, J. (1981), "The Present Status of the Innateness Hypothesis," in *RePresentations*, Cambridge (MA): MIT Press, pp257-316

Fodor, J. (1987), *Psychosemantics*, Cambridge (MA): MIT Press

Fodor, J. (1991), *A Theory of Content and Other Essays*

Fodor, J. (1998), *Concepts: Where Cognitive Science Went Wrong*, Cambridge (MA): MIT Press

Fodor, J. (2004), "Having Concepts: A Brief Refutation of the Twentieth Century," with "Reply to Commentators," *Mind and Language* 19(1): pp29-47, 99-112.

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

Greenberg, M. (2000), *Thoughts Without Masters: Incomplete Understanding and the Contents of Mind*, D. Phil dissertation, University of Oxford.

Horwich, P. (1998), *Meaning*, Oxford University Press

Jackendoff, R. (1983), *Semantics and Cognition*, Cambridge: MIT Press

Jackson, F. (1998), *From Metaphysics to Ethics: a Defense of Conceptual Analysis*, Oxford University Press

Kaplan, D. (1969), "Quantifying In," in *Words and Objections*, ed. by D. Davidson and J. Hintikka; Dordrecht: Reidel, pp178-214

Katz, J. (1972), *Semantic Theory*, New York: Harper & Row

Lewis, D. (1972), "How to Define Theoretical Terms," *Journal of Philosophy*, 67:427-44

Millikan, R. (2000), *On Clear and Confused Ideas*, Cambridge Univ Press

Morton, A. (1997), *A Guide Through the Theory of Knowledge*, Oxford: Blackwell

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

Peacocke, C. (1998), "Implicit Conceptions, the *A Priori* and Identity of Concepts," in *Concepts*, ed. By E. Villaneuva, vol 9 of *Philosophical Issues*, Atascadero: Ridgeview Press

Peacocke, C. (2004), *The Realm of Reason*, Oxford University Press

Pietroski, P. (2003), The Character of Natural Language Semantics, in *Epistemology of Language*, ed. By A. Barber, Oxford University Press, pp217-256.

Putnam, Hilary (1962/1975), "The Analytic and the Synthetic," in *Mind, Language and Reality: Philosophical Papers*, vol 2; Cambridge University Press, pp33-69

Putnam, H. (1965/75), "Dreaming and Depth Grammar," in *Mind, Language and Reality: Philosophical Papers*, vol 2; Cambridge University Press, pp304-24

Putnam, H. (1975), "The Meaning of `Meaning,'" *Mind, Language and Reality: Philosophical Papers*, vol 2; Cambridge University Press, pp215-71

Pylyshyn, Z. (2003), *Seeing and Visualizing: It's Not What You Think*, Cambridge (MA): MIT Press

Quine, W. (1953a), "On What There Is," in *From a Logical Point of View*, New York: Harper & Row, pp1-9

Quine, W. (1953b), "Two Dogmas," in *From a Logical Point of View*, New York: Harper & Row, pp20-46

Rey, G. (1981), "What Are Mental Images?", in *Readings in the Philosophy of Psychology*, vol. 2, ed. by N. Block, Cambridge: Harvard Univ. Press, 1981, pp 117- 1 27.

Rey, G. (1983/1999), "Concepts and Stereotypes," *Cognition* 15 (1983), pp237-262, re-printed in *Concepts: Core Readings*, ed. by E. Margolis and S. Laurence, MIT Press

Rey, G. (1985), *Cognition* 6. "Concepts and Conceptions," *Cognition* 19, pp297-303

Rey, G. (1997), *Contemporary Philosophy of Mind: a Contentiously Classical Approach*, Oxford: Blackwell

Rey, G. (1996), "Resisting Primitive Compulsions," contribution to symposium on C. Peacocke, *A Study of Concepts, Philosophy and Phenomenological Research*, Vol. LVI, #2, pp419-424

Rey, G. (1998), "What Implicit Conceptions Are Unlikely To Do," commentary on Christopher Peacocke, "Implicit Conceptions, Understanding and Rationality," *Philosophical Issues, 9: Concepts*, ed. by E. Villaneuvo; Atascadero: Ridgeview Press, pp89-92

Rey, G. (2003), "Intentionality and a Chomskyan Linguistics," in Barber, A. (ed.), *Epistemology of Linguistics*, Oxford Univ Press, pp140-86

Rey, G. (2004), "Fodor's Ingratitude and Changes of Heart?," Commentary on Fodor's, "Having Concepts: A Brief Refutation of the 20th. Century," *Mind and Language* 19(1), pp70-84.

Rey, G. (forthcoming-a), "Explanation, not Experience; Commentary on John Campbell's *Reference and Consciousness*", *Philosophical Studies*

Rey, G. (forthcoming-b), "The Rashness of Traditional Rationalism and Empiricism," *New Essays in Philosophy of Language and Mind*, Supplemental Volume, *Canadian Journal of Philosophy*, ed. by R. Stainton and C. Viger

Rey, G. (in prep), *Thinking about Nothing: Intentionality and Inexistence in a Chomskyan Theory of Mind*, Oxford University Press

Russell, B. (1912/59), *The Problems of Philosophy*, Oxford Univ Press

Slote, M. (1975), "Inapplicable Concepts," *Philosophical Studies*,
28:267-71

Taylor, K. (2003), *Reference and the Rational Mind*, CSLI Publications

Thomason, R. (2001), "Formalizing the Semantics of Derived Words," at
<http://www.eecs.umich.edu/~rthomaso/documents/nls/index.html>

Wittgenstein, L. (1953), *Philosophical Investigations*, New York: Macmillan

Index (starred entries that are especially important):

Brentano, F.

Burge, T.

Chomsky, N.

color

*concepts

concept of

*Classical View

*empty

*existential vs. purely intentional usage

*internal roles

one criterion

prototype view

(mental) content

asymmetric dependency theory of

*Externalist theories of

Euclid

*Externalism

Strong vs. Weak

*Fodor, J.

Horwich, P.

Jackson, F.

Katz, J.

Kripke, S.

Millikan, R.

modal judgments

phonetics

Palmer, S.

Peacocke, C.

properties

Putnam, H.

*Quine, W.

represent(ation)

existential vs. purely intentional usage

Rey, G.

Slote, M.

Wittgenstein, L.