

The Authority of Pragmatic Conceptual Analysis¹

Justin C. Fisher – Southern Methodist University
Working draft – please do not cite without permission

Abstract.

This paper defends *Pragmatic Conceptual Analysis*, a proposed empirical methodology for explicating philosophical concepts. This methodology attributes to our shared concepts whatever application conditions they would need to have in order best to continue delivering benefits in the ways they have regularly delivered benefits in the past. In the first stage of my argument I argue that Pragmatic Conceptual Analysis has what I call *normative authority*: we have practical and epistemic reason to adopt the explications that it delivers even if we think doing so requires stipulative revisions in the meanings of our concepts. I then use this normative authority to argue that Pragmatic Conceptual Analysis also has what I call *descriptive authority*: when we understand concept-meaning in the way we ought to understand it (in the way licensed by the normative authority of Pragmatic Conceptual Analysis) we see that, rather than being revisionary, Pragmatic Conceptual Analysis is a semantically conservative tool that uncovers (what we should think of as being) the meanings our concepts already have.

Introduction.

This paper explores the relative merits of different proposed methodologies for conceptual analysis. Any such methodology offers a way of determining, for any given concept,² a proposed *explication* for that concept – a proposed specification, in other terms, of what exactly that concept is applicable to. For example, *Intuitive Conceptual Analysis* is the familiar methodology³ that seeks explications which preserve the truth of intuitions surrounding the concepts being explicated.⁴ If reflection on possible cases reveals that we'd be intuitively willing to count all and only unmarried men as bachelors, then Intuitive Conceptual Analysis would explicate BACHELOR as unmarried man.

¹ ### [Acknowledgements removed for blind review.]

² I think of concepts as mental particulars that play a folder-like information-coordinating role in cognition. In this, I follow a rich tradition in cognitive psychology and the philosophy of cognitive science. For a good introduction to this tradition see Lawrence & Margolis (1999), and for a view of concepts very similar to my own see Millikan (1998, 2000). This tradition may be contrasted against an equally rich philosophical tradition that takes concepts to be abstract entities.

³ Like other methodologies we'll consider, Intuitive Conceptual Analysis is best thought of as a general methodology with various particular versions. Different versions might disagree about whose intuitions we should consult (ordinary folk or experts?) and about how we should resolve trade-offs between conflicting intuitions.

⁴ One advocate of Intuitive Conceptual Analysis is Frank Jackson, who writes:

[H]ow should we identify our ordinary conception? The only possible answer, I think, is by appeal to what seems to us to be most obvious and central about free action, determinism, belief, or whatever, as revealed by our intuitions about possible cases. [...] Intuitions about how various cases, including various merely possible cases, are correctly described in terms of free action [etc...] are precisely what reveal our ordinary conceptions of free action [etc...] (Jackson 1998, p. 31).

Other methodologies prioritize other factors over intuitions. For example, the methodology I call *Naturalized Conceptual Analysis* seeks explications that are true of whatever natural kind has been playing an appropriate role in causing us to use the given concept.⁵ If the samples that have caused us to form our concept WATER and to associate information with it were, in fact, all samples of the natural kind H₂O, then Naturalized Conceptual Analysis would explicate WATER as H₂O regardless of whether there were pre-existing intuitive links between water and hydrogen.

Alternatively, the methodology I call *Pragmatic Conceptual Analysis* seeks explications that would preserve the ways in which a shared concept has regularly been beneficially used.⁶ If, as Johnston (1912, Chapter II) recounts, the Huron Indians employed a concept whose primary beneficial use was to recognize and treat scurvy, then Pragmatic Conceptual Analysis would explicate this concept as meaning scurvy, regardless of whether the Huron intuitively thought this ailment was essentially a curse by enemy spell-casters, and regardless of whether they frequently misdiagnosed a variety of other ailments as falling under that concept.

(Readers sometimes balk at my labeling the latter two methodologies as forms of “Conceptual Analysis”. I use this label because these methodologies share the general goal of traditional conceptual analysis: carefully exploring the relevant features of pre-existing concepts to state the application conditions of those concepts in other terms. Admittedly, these methodologies don’t involve “analysis” in the old-fashioned sense of breaking concepts apart into component parts. But virtually everyone is now convinced by psychological research⁷ that few, if any, concepts have the neat definitional structure that would be required for them to be “broken apart” in this way. So, even contemporary practitioners of Intuitive Conceptual Analysis think “analysis” doesn’t involve breaking concepts apart into component parts, but instead examining the roles concepts play in broader cognition. Naturalized and Pragmatic Conceptual Analysis hold that the scope of analysis is a bit more broad – including the roles concepts play in interaction with external objects – but that doesn’t make this examination any less of an analysis.)

⁵ Naturalized Conceptual Analysis enjoys a significant following, including proposed analyses of KNOWLEDGE (Quine 1969, Kornblith 2002), EMOTION (Griffiths 1997), COLOR (e.g., Hilbert 1987), CONSCIOUSNESS (Dennett 1991), and MORAL GOODNESS (e.g., Boyd 1988). Kornblith offers an especially clear characterization of this methodology:

We begin, often enough, with obvious cases, even if we do not yet understand what provides the theoretical unity to the kind we wish to examine. *Understanding what that theoretical unity is* is the object of our study, and it is to be found by careful examination of the phenomenon, that is, something outside of us, not our concept of the phenomenon, something inside of us. (Kornblith 2002, pp. 10-11.)

Some readers are reluctant to concede the label ‘naturalized’ to this particular methodology, when other methodologies (including my own) might equally draw upon results of natural science. This concern is apt, but I know of no other plausible label with similar currency in the literature, so I reluctantly use this one.

⁶ Pragmatic Conceptual Analysis will be more carefully articulated in Section 1.1 below. For a more detailed development of Pragmatic Conceptual Analysis, including discussion of many different particular varieties of this general methodology, see (###). Closely related methodologies have been proposed in epistemology by Edward Craig (1990) and Sally Haslanger (1999), and in philosophy of science by James Woodward (2004). Haslanger offers an especially nice characterization:

[T]he best way of going about a project of normative epistemology is first to consider what the point is in having a concept of knowledge: what work does it, or (better) could it, do for us? And second, to consider what concept would best accomplish this work. (Haslanger 1999, pg 467)

⁷ For an overview of relevant research, see Laurence & Margolis (1999).

We'll say much more about these three methodologies (and others) in what follows. But first, let's consider some virtues we might hope proposed methodologies like these would display.

Let us say that a proposed methodology has *normative authority* if and only if, for any given concept, that methodology offers an explication that the people employing that concept would *have good reason to* adopt. For example, the narrator of *Moby Dick* insists that his concept FISH is correctly applicable to whales (“I take the good old fashioned ground that the whale is a fish, and call upon holy Jonah to back me”) but he also acknowledges that whales differ in significant ways from most fish: “in brief they are these: lungs and warm blood; whereas all other fish are lungless and cold blooded” (Melville 1851, pg 132). Regardless of whether you agree with the narrator's thought that his FISH concept could correctly be applied to whales, you might think he would *have good reason to stipulate* that henceforward he will use his FISH concept as biologists do, reserving this concept for those scaly lungless creatures who share many more scientifically interesting traits with one another than they do with whales. Some of these good reasons might be *practical* in that they would help one achieve beneficial ends, like efficiently predicting which creatures will be most likely to capsize one's boat. And some good reasons might be *epistemic* in that they would help one achieve a well-systematized catalogue of true beliefs about the world. If the narrator and his compatriots would have good practical and epistemic reasons to stipulate this narrower meaning for FISH, then a normatively authoritative methodology would recommend this explication for them.

There's a clear sense in which normative authority is the most important virtue we could hope for in a methodology for conceptual analysis. Once we know what proposed explications we have good reason to adopt, that tells us what we usually most want to know, and there's usually little use for other proposed explications. Suppose methodology N has normative authority while methodology D does not. Since N has normative authority, we'll have reason to embrace the explications N delivers. Not so for D. Hence, whenever we have reason to prefer one methodology's explications over the other's, we'll have reason to adopt N's. So, whatever other virtues D might have, these virtues must be quite academic indeed, as these virtues can easily be trumped by our reasons to adopt N's explications instead.

Despite the clear importance of normative authority, many people have instead yearned for another virtue which I call *descriptive authority*. A proposed methodology has *descriptive authority* if and only if, for any given concept, that methodology tells us what things, both actual and possible, that concept *actually is correctly applicable to*. So (given common presumptions about correct applicability), a descriptively authoritative methodology would need to tell us that our concept WATER is correctly applicable to all and only the instances of H₂O and that SQUARE is correctly applicable to all and only the regular four-sided planar figures.

I will often abbreviate talk of “correct application conditions” to talk of “meaning.” So I say that WATER *means* H₂O and that descriptively authoritative methodologies are methodologies that correctly capture the *meanings* of all concepts. It's worth emphasizing that this notion of *concept-meaning* is purely extensional across possible worlds. So, for example, it ignores any ‘intensional’ differences between TRIANGLE and TRILATERAL or between WATER and H₂O; and instead takes each pair to be alike in ‘meaning’.

Getting an extensionally adequate understanding of our concepts will be challenge enough for this paper. After we resolve that challenge, you might hope to go on and draw finer-grained or ‘intensional’ distinctions between meanings.

It’s controversial which methodologies have descriptive authority because it’s controversial which theory of concept-meaning is correct. Each of the three methodologies mentioned above closely corresponds to its own class of semantic theories in philosophy of mind and language (summarized in Table 1). Intuitive Conceptual Analysis corresponds to Descriptivist theories which take a concept’s meaning to be determined by the descriptions we tacitly associate with the concept, as our intuitions plausibly would reflect such inner descriptions.⁸ Naturalized Conceptual Analysis corresponds to Causal or Informational semantic theories that take a concept’s meaning to be determined by what natural kinds have played an appropriate role in causing our usage of that concept.⁹ And Pragmatic Conceptual Analysis corresponds to pragmatic or teleosemantic theories that take a concept’s meaning to be determined by the ways in which that concept has beneficially been used.¹⁰ Whichever one (if any) of these semantic theories turns out to be correct, it’s likely that the corresponding methodology will be descriptively authoritative.

Semantic Theory	Methodology for Conceptual Analysis
<p>Descriptivist theories hold that a concept is correctly applicable to whatever fits the description(s) we have tacitly associated with that concept.</p>	<p>Intuitive Conceptual Analysis seeks an explication that best fits our intuitions involving a concept (which hopefully will reflect the descriptions hidden inside us).</p>
<p>Causal/Informational theories hold that a concept is correctly applicable to whatever has played an appropriate role in causing its usage.</p>	<p>Naturalized Conceptual Analysis identifies some clear instances of a concept (typically instances that caused us to use that concept) and then seeks a natural kind that includes those instances.</p>
<p>Teleosemantic/Pragmatic theories hold that a concept is correctly applicable to whatever would best sustain existing patterns in its beneficial usage.</p>	<p>Pragmatic Conceptual Analysis seeks an explication that will best preserve the patterns of beneficial usage for a given concept.</p>

Table 1. Three package deals (rows) each combine a semantic theory and a methodology for conceptual analysis.

Interestingly, the left column of Table 1 lists semantic theories in approximate *increasing* order of popularity in contemporary philosophy of mind and language. Most naturalized semanticists have abandoned Descriptivism, and moved on to either Causal/Information or Teleo/Pragmatic theories which instead take meaning to be determined in large part by patterns of causal interaction between our concepts and things in the world. In contrast,

⁸ Descriptivists include Frege (1892), Russell (1905), and Strawson (1950).

⁹ Causal/informational theorists include Kripke (1972), Evans (1973), Boyd (1988), Fodor (1990), and Rupert (1999). Kornblith (2002, pp. 12-13) notes the affinity between Naturalized Conceptual Analysis and causal theories of meaning.

¹⁰ Teleo/pragmatic theorists include James (1906), Ramsey (1927), Millikan (1984), Appiah (1986), Papineau (1987), Dretske (1988), Whyte (1990), and Blackburn (2005).

the right-hand column lists philosophical methodologies in *decreasing* order of current popularity. Intuitive Conceptual Analysis is still the main stock-in-trade of analytic philosophy; a minority of empirically minded philosophers have moved on to Naturalized approaches; while very few analytic philosophers defend methodologies that would fit closely with teleo-semantics.¹¹ Indeed, it is this conspicuous vacancy that Pragmatic Conceptual Analysis is intended to fill. This is an unstable situation. Sooner or later some sort of correction will need to be made to resolve it. Either (a) philosophers of mind and language will need to reincarnate descriptivism,¹² or (b) we'll need to abandon our pretense to descriptive authority, or else (c) we'll all need to adopt methodologies that better fit our state-of-the-art semantic theories.

Since it's controversial which methodologies have descriptive authority, it's also controversial what the relation is between descriptive and normative authority. An optimist might hope that most of our concepts already do mean what we would have reason to stipulate they mean. If so, then there won't be much difference between a *descriptively* authoritative methodology and a *normatively* authoritative one. However, until we establish that such optimism is warranted,¹³ it is reasonable to remain agnostic and allow that some (or even many) of our concepts might currently be suboptimal, actually meaning something other than what we would have reason to stipulate they mean.

The goal of this paper is to show that optimism is warranted. I will offer reasons to think that Pragmatic Conceptual Analysis is both normatively and descriptively authoritative. We will intentionally develop Pragmatic Conceptual Analysis in a way that ensures that it has *practical* normative authority (in Section 1.1). It will then be fairly straightforward to argue that it also *epistemic* normative authority (in Section 1.2).

The question of whether Pragmatic Conceptual Analysis also has descriptive authority will be much trickier. This question is equivalent to asking whether Pragmatic Conceptual Analysis reveals the correct application conditions or 'meanings' of our concepts. As noted above, it's controversial how 'concept-meaning' should be understood, so it's

¹¹ I listed the few clear examples I am aware of in note 6. Some past methodologies share much of the spirit of Pragmatic Conceptual Analysis, including James (1906) and Carnap (1950, pg 7). Other contemporary authors, including Schmidtz (2006, pp. 11-12) and Jackson (1998, p. 45) have proposed something like Pragmatic Conceptual Analysis as a fallback strategy to employ when Intuitive Conceptual Analysis is inconclusive or yields undesirable results.

¹² Some contemporary philosophers aim to do just that, including Lewis (1984) and Chalmers (2002, 2004).

¹³ One potential ground for optimism is the common belief that semantic theories should obey a *Principle of Charity* which prohibits being needlessly uncharitable in our attributions of meaning. See Quine (1960), Lewis (1983), Davidson (1984), Dennett (1987), and Henderson (1994). It is usually uncharitable to attribute meanings to someone that she would have good reason to stipulatively change. E.g., it would be uncharitable to insist that the Huron concept meant something other than scurvy, when it's clear they had been using – and had every reason to continue using – this concept as their means for recognizing and curing scurvy. Hence a principle of charity demands that any *descriptively* authoritative methodology must also (at least usually) be *normatively* authoritative. Combining this observation with my argument (in sections 1.1–1.2) that Pragmatic Conceptual Analysis has a unique claim to normative authority would yield a charity-based argument for concluding that any descriptively authoritative methodology must be much like Pragmatic Conceptual Analysis. If you already accept a principle of charity, this charity-based argument should sway you toward Pragmatic Conceptual Analysis. But it's controversial whether we should accept a principle of charity, so I won't rest the weight of my main argument on this controversial premise. (Instead I think a version of the Principle of Charity can be justified as a corollary of my bootstrapping argument.)

unclear what exactly this question is asking. To approach a tricky question like this, we should first make clear what's at issue. This suggests the following three-stage strategy:

Stage 1. Seeking Normative Authority. First, we'll seek a normatively authoritative methodology that we will have practical and epistemic reason to use to clarify our tricky question.

Stage 2. Explicating 'Concept-Meaning'. Second, we'll use the methodology we identified in Stage 1 to explicate 'concept-meaning' thereby clarifying the tricky question we started out with.

Stage 3. Assessing Descriptive Authority. And third, we'll answer the question thus clarified: does Pragmatic Conceptual Analysis reveal concept-meaning, as we came to understand it in Stage 2?

We'll consider Pragmatic Conceptual Analysis when we get to Stage 3 and ask whether it reveals concept-meaning, appropriately understood. But Pragmatic Conceptual Analysis will arise in the earlier stages too. For I will argue in Stage 1 that we generally have good reason to use Pragmatic Conceptual Analysis to explicate our concepts. So, in Stage 1 we'll discuss our reasons for using Pragmatic Conceptual Analysis; in Stage 2 we'll use Pragmatic Conceptual Analysis to explicate 'concept-meaning'; and in Stage 3 we'll ask whether the resulting understanding of concept-meaning is the sort of meaning that Pragmatic Conceptual Analysis reveals for all concepts. I call this a 'bootstrapping argument' because it allows us legitimately to use Pragmatic Conceptual Analysis to help us draw conclusions about itself.

The answer to our tricky question will end up being a qualified 'yes': Pragmatic Conceptual Analysis does reveal concept-meaning, *appropriately understood*. In the final section I argue that this italicized qualifier is less significant than it might seem; for most purposes we may dispense with it entirely.

Stage 1. Normative Authority

1.1. Practical Normative Authority

In Stage 1 of our bootstrapping argument, our goal is to identify a methodology with *normative authority*: a methodology whose explications we'll have practical and epistemic reason to adopt regardless of whether we think of those explications as reflecting the existing meanings of our concepts or as revising their meanings in useful ways. This normative authority will license our using this methodology in Stage 2 to explicate 'concept-meaning' and thereby to clarify the tricky question of which methodology reveals concept-meaning, a question we'll then take up in Stage 3.

Let's begin by considering the case of free action. Frank Jackson, a staunch advocate of Intuitive Conceptual Analysis, consults his intuitions and concludes:

Speaking for my part, my pre-analytic conception of free action is one that clashes with determinism. (Jackson 1998, p. 44).

Jackson himself notices drawbacks of his incompatibilist explication of FREE ACTION. Given that determinism is (near enough to) true, Jackson's explication would force us to say that *no* actions are free. This would undercut many useful distinctions we ordinarily use this concept to make, like the distinction between the “free” acts of ordinary vandals and the “unfree” acts of sleep-walkers. Faced with these failings of Intuitive Conceptual Analysis, Jackson proposes a fallback strategy:

It is, thus, only sensible to seek a different but ‘nearby’ conception that does, or does near enough, the job we give [to the concept being analyzed] in governing what we care about, our personal relations, our social institutions of reward and punishment, and the like, and which is realized in the world. (Jackson 1998, pg 45)

Jackson's fallback strategy is a passable characterization of Pragmatic Conceptual Analysis. The crucial difference between Pragmatic Conceptual Analysis and Intuitive Conceptual Analysis is that, where Intuitive Conceptual Analysis would follow intuitions surrounding a concept into a counter-productive explication of that concept, Pragmatic Conceptual Analysis would instead follow the patterns of beneficial usage surrounding that concept to yield an explication that preserves those beneficial uses, and hence is practically worth adopting. It's gratifying that Jackson, a key advocate of a competing methodology, admits that, when his methodology and mine disagree, it is “only sensible” to follow the guidance of mine.¹⁴ (With enemies like this, who needs friends?)

We can work backward from the goal of identifying a methodology with practical normative authority to spell out, in much more detail than Jackson's passable characterization, the sort of Pragmatic Conceptual Analysis that would best meet this goal. We're looking for a methodology whose explications we'll have practical reason to employ because of the expected beneficial effects of employing those explications as our standard for when to apply the explicated concepts.

This paper stakes no particular claim regarding which ‘benefits’ we have practical reason to pursue. I presume that there are some such benefits – e.g., typically, the benefit of getting food when one is hungry. Even without a general account we can often identify some clearly pursuit-worthy benefits in particular cases: curing diseases, preventing violent crime, helping us avoid danger. Such uncontroversial benefits will be enough to sustain the conclusions of this paper. For the purposes of this paper, I remain neutral regarding whether pursuit-worthy benefits involve survival, health, happiness, desire satisfaction, flourishing, and/or something else. I also remain neutral regarding whether all pursuit-worthy benefits are self-involving, or whether we sometimes have reason to bring about effects that directly involve only other people. I leave it to other philosophers to provide a general account of what exactly all the other pursuit-worthy benefits are.¹⁵

If we're going to have practical reason to employ a methodology's explications, the employment of these explications will need to be correlated with beneficial effects. This could happen in two ways: (1) the methodology could be explicitly spelled out in terms of

¹⁴ Jackson would insist that this “sensible” application of Pragmatic Conceptual Analysis forces us to revise the meaning of FREE ACTION. In Stage 3, I will argue against this claim. All I need in Stage 1 is that we have reason to adopt the explications Pragmatic Conceptual Analysis delivers, and Jackson accepts this.

¹⁵ For further discussion of the different understandings of “benefit” that might be useful for different purposes in Pragmatic Conceptual Analysis, see (###).

beneficial effects, or (2) it could be spelled out in terms of other factors that happen to be correlated with beneficial effects. Option (2) risks falling short of our goal and offers no significant advantages over option (1). It risks falling short of our goal because whatever factors it mentions might turn out not to be so well correlated with practical benefits after all (as for example, we've seen that intuitions needn't be correlated with benefits). Option (2) offers no real advantages because no matter what factors it mentions (be they intuitions, causal baptisms, associated metaphors, or whatever), these factors *can't* be better correlated with practical benefits than practical benefits are with themselves. So, whatever version of option (2) you might want to adopt, there's a version of option (1) that will do at least as well: one that specifies the relevant benefits directly, rather than coyly talking about some correlate. Hence, I will consider only versions of option (1) – i.e., only methodologies that are spelled out in terms of practical benefits.

If we're going to have practical reason to adopt a proposed explication *as an explication of an existing concept*, rather than to coin some new concept, then, as Jackson said, the proposed explication must somehow be 'nearby' the original concept. However, the relevant notion of 'nearness' involves preserving beneficial uses rather than intuitions. Giving up old intuitions needn't be practically inadvisable – indeed, as Jackson noted, old intuitions often conflict with the explications that are most practically advisable. In contrast, if a proposed explication were to give up on very much of the *beneficialness* of the original concept, then it would be practically advisable *both* to keep the original for its benefits *and* to coin a new concept with the explication as its stipulated meaning (to get its benefits as well). This means that a methodology with practical normative authority must propose explications that preserve enough of our concepts' existing patterns of beneficial usage that these proposals will be worth accepting *as explications*.

I noted above¹⁶ that I think of each individual concept as playing a folder-like role in someone's cognition. However, we'll typically have practical reason to preserve not just the beneficial uses of our own concepts, but also the beneficial uses of the corresponding concepts possessed by our compatriots. Let us define a *shared concept* as being composed of the various individual concepts held by compatriots who are disposed to pass back and forth the information they associate with these concepts, and, in cases of conflict, to debate which information should be associated with them.¹⁷ Since you typically have practical reason to preserve *any* regular beneficial use of a shared concept (not just those uses you happen to have already enjoyed yourself), a normatively authoritative version of Pragmatic Conceptual Analysis should be defined in terms of *shared* concepts rather than *individual*

¹⁶ See note 2.

¹⁷ One might think of shared concepts as *aggregations* of individual concepts, as *types* of individual concepts, or as *sets* of individual concepts. However, I prefer to think of shared concepts as *distributed historical individuals* in their own right, roughly on the model of biological species as distributed historical individuals (Hull 1978). Like biological species, shared concepts are restricted to relatively small portions of space-time; many of their population-level features are stably maintained by homeostatic processes; their members are arranged in lineages which display heritable variation; and they display cumulative change over time including a sort of adaptation to surrounding conditions. I find this understanding preferable because it highlights the sorts of dynamical coordination and interaction that are essential to shared concepts, and that are not displayed by garden variety aggregations, types, or sets. However, everything I say below might easily be reconstrued in terms of these more garden-variety entities.

concepts.¹⁸ (In cases where this distinction is not especially important, I will often use the term ‘concept’ for short.)

If we are to be confident that a proposed explication will continue to yield benefits much as the original shared concept did, then this explication can’t be based upon mere spurious, arbitrary, or gerrymandered patterns in the history of usage of the original concept. Instead this explication must be based upon robust, explanatorily relevant patterns of beneficial usage that are projectible into the future. This means that a methodology with practical normative authority must have much in common with an account of *good explanation*: it must detect which recurrent patterns of usage *explain* how selective application of a shared concept has been beneficial, and this explanation must be *simple, robust* and *natural* enough that it can be expected to continue to apply in the future.

This brings us quite a long ways towards articulating a normatively authoritative version of Pragmatic Conceptual Analysis. Such a methodology would involve two steps:

Step 1: First, we seek a ‘job description’ for the given shared concept: a good (simple, robust, natural) explanation of how our selective application of that concept to some things but not others has regularly led to beneficial effects – i.e., effects that the individual employing the concept had practical reason to pursue.

Step 2: With such a job description in hand, we then seek to identify an explication of that concept, which, if employed as the standard for applying the concept, would do optimally well at yielding beneficial effects in accord with this job description.

In some cases this methodology will lead straightforwardly to quite determinate explications. Suppose, as seems quite plausible, that a good explanation for how our usage of the concept WATER has been beneficial would say that this concept has regularly been used to predict which things (e.g., puddles but not rocks) will display one or more of the many characteristic features of H₂O. In order to do this job optimally well, we would need to apply WATER to all and only the instances of H₂O. Or suppose, perhaps less plausibly, it turns out that the correct job description for FREE ACTION includes only this concept’s role in distinguishing people who have well-functioning deliberative systems that are responsive to rewards and punishment from people whose actions would instead be better influenced by other means like medical treatment or isolation. This finding would give us practical reason to accept a compatibilist explication of FREE ACTION as something like “action that is susceptible to control by a well-functioning deliberative system.”

Unfortunately, in many cases – including, perhaps, the case of free action – things will be more complicated. The job descriptions for many concepts will likely include multiple slightly incompatible strands of beneficial usage, and we will be faced with a choice between various explications, none of which fully preserves all strands. The case of JADE

¹⁸ This move to shared concepts accommodates something much like Putnam’s (1973) “linguistic division of labor.” However, Putnam imagined his division as involving a few experts who determine meanings for everyone else, whereas Pragmatic Conceptual Analysis allows that the determination of meaning might be distributed more broadly across many people with varying degrees of (in)expertise. So long as a community together provide a large enough pool of beneficial usage, they can collectively determine the meaning of a shared concept, even if no individual in the community is experienced enough to fully do so herself.

is one famous example, where one strand of usage involved jadeite and another nephrite. Different versions of Pragmatic Conceptual Analysis might resolve these tradeoffs in different ways: e.g., by somehow picking one important strand to optimize, or by ‘satisficing’ to maintain multiple strands of beneficial usage somewhat well. (This is perfectly analogous to the ways in which different versions of Intuitive Conceptual Analysis might resolve conflicting intuitions.)

However, for present purposes, these complexities won’t matter. The differences between these versions of Pragmatic Conceptual Analysis will usually be fairly slight. Furthermore, our practical and epistemic purposes will usually be similarly well-served by adopting *any* clear explication that preserves a maximal portion of our original beneficial usage (and by coining new concepts, where necessary, to handle whatever potential benefits get left out). E.g., for JADE we chose to satisfice and include both jadeite and nephrite, whereas for FISH we chose to keep just the ichthians and exclude the whales. But things would have worked out similarly well had we instead reserved JADE for jadeite and coined a new concept (like NEPHRITE) for nephrite, or had we used FISH for all aquatic vertebrates including whales, and started employing a different concept (like ICHTHIAN) for just the ichthians. Often, multiple ways of resolving such conflicts are similarly practically advisable. So long as we’re employing a version of Pragmatic Conceptual Analysis that chooses *some* practically advisable way, we’ll have practical reason to employ it.

We saw above that such practical reasons favor Pragmatic Conceptual Analysis over Intuitive Conceptual Analysis: we have practical reason to choose explications that preserve patterns of beneficial usage even at the cost of denying folk intuitions. Parallel considerations arise against *any* methodology willing to sacrifice beneficial-ness in order to preserve independent features of our pre-existing concepts (surrounding intuitions, causal baptisms, associated metaphors, or whatever). Unlike such competitors, Pragmatic Conceptual Analysis is explicitly geared towards delivering explications that preserve the beneficial uses of our concepts. When we combine this feature of Pragmatic Conceptual Analysis with the premise that we have practical reason to pursue such benefits, we get the conclusion that we have instrumental practical reason to embrace the explications Pragmatic Conceptual Analysis delivers, as a way to achieve these benefits.

The normative authority of Pragmatic Conceptual Analysis comes with a few caveats. Sometimes we know that old ways of achieving benefits will no longer work and sometimes we know of promising new ways to achieve benefits.¹⁹ Most good guides are

¹⁹ E.g., a pragmatic conceptual analysis of PHLOGISTON might tell us how one could best salvage the relatively limited benefits that phlogiston theory delivered, e.g., perhaps by counting all and only oxidizable materials as instances of phlogiston. This might be a useful and charitable understanding of what phlogiston theorists were on about; but we might nevertheless have overriding reason to stop wasting time on this theory and instead to pursue the novel benefits afforded by chemistry. Many other cases (e.g., WITCH or CALORIC) that sometimes have been deemed “non-referring concepts” succumb to a similar treatment: Pragmatic Conceptual Analysis provides a charitable explication that retains what little work the concept has been doing, yet this work often is not important enough to merit keeping the concept. I give a full treatment of issues surrounding so-called “non-referring concepts” in (###) but it is outside the scope of the present paper to rehearse that here. In the end, it matters little what explication we give for degenerate concepts that largely fail to track anything interesting at all, especially when we’re likely to abandon such concepts anyway. Hence, such cases can provide no strong reason to favor any proposed methodology over another.

defeasible in this way. It's reasonable to disregard the advice of even the best trail-guides when we know the trail has recently changed or when we know of potential benefits our guide hasn't considered. Nevertheless, good trail-guides have normative authority – we generally have reason to heed their advice. Similarly, even though Pragmatic Conceptual Analysis is defeasible, it too has normative authority. So long as we remain alert for competing reasons and potential defeaters, we generally have reason to adopt the explications that Pragmatic Conceptual Analysis delivers.²⁰

1.2. Epistemic Normative Authority

The above reasons for following Pragmatic Conceptual Analysis are *practical reasons* – they involve bringing about practical consequences worth pursuing. There is some precedent for highlighting practical reasons to adopt philosophical conclusions: Pascal's famous wager highlights a practical reason to believe that God exists (Pascal 1910) and various Dutch Book arguments highlight practical reasons for embracing probabilism (Joyce 1998). Still, one might doubt the value of such arguments. If an eccentric billionaire were to offer tempting financial incentives to adopt, say, a causal theory of concept-meaning, we would rightly deem this irrelevant to philosophical questions regarding concept-meaning.²¹ Why should the practical incentives Pragmatic Conceptual Analysis offers be any more philosophically relevant?

For readers who would like something more than mere practical reasons for adopting my philosophical conclusions, I will now argue that, closely linked to our practical reasons for following Pragmatic Conceptual Analysis, we also have *epistemic reasons* to follow this methodology. My argument will take the following form:

- P1.** We have epistemic reason to follow any conceptual maintenance strategy that consistently enables our epistemic successes.
 - P2.** Positive reinforcement learning of categorization behavior is one such strategy.
 - P3.** Pragmatic Conceptual Analysis is effectively the same strategy as positive reinforcement learning of categorization behavior.
-

C. Hence, we have epistemic reason to follow Pragmatic Conceptual Analysis.

I will work with a traditional understanding of epistemic success as believing interesting truths without believing falsehoods, though other understandings might work just as well. Just as we have *practical reason* to follow strategies that consistently enable us to achieve

²⁰ One might hope to describe a version of Pragmatic Conceptual Analysis that spells out what to do even in extraordinary cases where these caveats apply. However the extraordinary cases are so diverse that I fear very little can be said in general about what to do in them, short of extremely general platitudes about practical reasoning like “weigh the costs and benefits of different proposals and choose one with maximal expected utility”. My own preference is not to include such open-ended platitudes in a proposed methodology for conceptual analysis. Instead, I prefer to describe a less open-ended methodology that sticks quite closely to established patterns of beneficial usage and leads straightforwardly to (fairly) determinate explications which we'll normally have defeasible practical reason to adopt. But this lack of open-endedness forces me to concede that, in certain extraordinary circumstances, other factors may defeat these practical reasons, and call for *ad hoc* general purpose practical reasoning instead.

²¹ Thanks to David Chalmers for this example.

practical success, we have *epistemic reason* to follow strategies that consistently enable us to achieve *epistemic* success.

This consequentialist construal of ‘epistemic reasons’ may strike some readers as being quite broad. I happily allow that other construals might also be useful, e.g. ones that tie ‘epistemic reasons’ more directly to having evidence that supports a conclusion. However, I want to emphasize that, so long as we identify appropriate epistemic goals as ‘epistemic successes’, my proposed understanding of ‘epistemic reasons’ will reflect the strategies we regularly use to achieve these appropriate epistemic goals. Hence there’s a good sense in which my ‘epistemic reasons’ reflect *good epistemic practices* even if they arguably don’t reflect a traditional understanding of ‘epistemic reasons’. If I can show that *good epistemic practices* would lead people to embrace my conclusions, that will be good enough for me.

It is an empirical question which strategies underlie our regular epistemic successes. Many robust findings of empirical psychology involve *positive reinforcement learning*. This idea was formulated in Edward Thorndike’s (1898, 1911) ‘law of effect’:

Of several responses made to the same situation, those which are accompanied or closely followed by satisfaction to the animal will, other things being equal, be more firmly connected with the situation, so that, when it recurs, they will be more likely to recur. (Thorndike, 1911, p. 244)

Positive reinforcement learning occurs not just for people’s overt behaviors but also for their application of concepts. For example, Ido Erev and colleagues have shown that a model based on Thorndike’s law of effect does better than competing models at accounting for how human subjects learn to perform binary categorization tasks like the task of guessing someone’s gender on the basis of their height (Erev 1998).

Mechanisms of positive reinforcement also effectively operate at the *inter*-personal level. We attempt to reproduce not just our own practical successes but also the practical successes of others. We tend to learn how to use shared concepts from the people who have demonstrated that they can use these concepts successfully, and we tend to try to emulate the practical successes that these people have enjoyed and not their failures. (For a review of relevant experimental evidence, see Henrich & McElreath, 2001.)

Positive reinforcement learning of categorization behavior is directly conducive to practical success, but it is indirectly conducive to epistemic success as well. For it’s generally hard to achieve consistent practical success without having true beliefs about interesting matters – i.e., without having epistemic success as well. When we learn to continue applying our concepts in ways that yield practical success, this is an effective strategy for focusing our concepts upon referents regarding which we can consistently achieve epistemic success too. Hence positive reinforcement learning of categorization behavior is one of our good epistemic practices that we have epistemic reason to continue.

Positive reinforcement learning of categorization behavior is effectively the same strategy as Pragmatic Conceptual Analysis: both seek to discover which ways of applying a concept have regularly been beneficial and then adopt these as our ways of applying that concept in the future. Pragmatic Conceptual Analysis is a self-conscious and deliberate way of doing something we already do subconsciously and automatically all the time.

I have argued (P1) that we have epistemic reason to follow whichever strategies underlie our regular epistemic successes. I then offered empirical reasons for thinking (P2) that positive reinforcement learning of categorization behavior is one of these tried-and-true epistemic strategies. Since (P3) Pragmatic Conceptual Analysis is just a formal articulation of this tried-and-true epistemic strategy, this yields the conclusion (C) that we have epistemic reason to follow Pragmatic Conceptual Analysis. Pragmatic Conceptual Analysis has not just practical normative authority but *epistemic normative authority* as well.²²

It is worth noting the close connection between practical reasons and epistemic reasons. Our epistemic reasons are grounded in our tried-and-true epistemic strategies, and an epistemic strategy wouldn't have survived to be tried-and-true if it weren't conducive to practical success. Conversely, our practical strategies often rely upon the accuracy of our beliefs, so our practical reasons are often geared towards achieving epistemic success. Given this close connection, it's unlikely that we would have strong *epistemic* reasons to employ a conceptual-maintenance methodology that we didn't also have strong *practical* reasons to employ. This, together with our (Section 1.1) conclusion that *practical* reasons strongly favor Pragmatic Conceptual Analysis over competing methodologies, entails that it's also unlikely that any competing methodology will be able to claim the strength of *epistemic* normative authority that I have claimed for Pragmatic Conceptual Analysis. And this means that competing methodologies would likely be missing a crucial premise, were they to try to concoct their own bootstrapping arguments analogous to mine.²³

Stage 2. Explicating CONCEPT-MEANING

2.1. Taking Stock.

We've now completed Stage 1 of the bootstrapping argument by establishing that Pragmatic Conceptual Analysis has both practical and epistemic normative authority. Before proceeding to Stage 2, it is worth emphasizing the significance of this accomplishment.

One way of seeing this significance is by considering a worry often voiced by non-philosophers and even by many philosophers. The worry is that analytic philosophy is

²² You might be inclined to think that epistemic norms must be semantically conservative. Combining this premise with my conclusion that epistemic norms endorse Pragmatic Conceptual Analysis would yield the further conclusion that Pragmatic Conceptual Analysis is semantically conservative. I think this further conclusion is correct but it shouldn't be this easy to establish. For I don't think we can help ourselves to the premise that all epistemic norms are semantically conservative. Normal successful epistemic practice might sometimes require that we make stipulative semantic revisions, e.g., to eliminate problematic sorts of semantic indeterminacy. Since *some* broadly epistemic norms plausibly call for semantic revision, we can't yet rule out the possibility that Pragmatic Conceptual Analysis might be semantically revisionary as well. It will take more argument to establish that this methodology is semantically conservative.

²³ This also provides an answer to an objection readers sometimes raise: what's to stop us from construing "benefits" in a such way that Pragmatic Conceptual Analysis would collapse into one of the other methodologies? E.g., if we define the "beneficial effects" of a concept as those that maximally satisfy our intuitions surrounding that concept, wouldn't the resulting version of Pragmatic Conceptual Analysis be equivalent to Intuitive Conceptual Analysis? However, if we were to adopt some non-standard understanding of "benefits effects" like this, then (presumably) we won't have practical or epistemic reason to focus upon pursuing such effects, so the resulting bastardized version of Pragmatic Conceptual Analysis would lack normative authority. For Pragmatic Conceptual Analysis to have normative authority, we must understand "benefits" as something we actually do have practical reason to pursue (as I stipulated in section 1.1).

utterly academic: philosophers spend entire careers with their heads in the clouds, pumping intuitions about fanciful cases and debating minutiae that have no bearing on real life. Pragmatic Conceptual Analysis avoids this worry. Pragmatic Conceptual Analysis demands that philosophers stay down to earth, looking at the beneficial work that our concepts have been doing for us, and seeking explications that will allow them to do this work more consistently. When done well, this sort of analytic philosophy can make a positive difference in people's lives and in our understanding of the world. We now have a fairly clear proposal regarding how we might fruitfully approach a wide variety of philosophical problems and we have good practical and epistemic reasons to approach problems in this way. This is the most important virtue that we could hope for in a philosophical methodology – anything more would be icing on the cake.

In light of all this, one might ask: *If we now know how to figure out what meanings we ought to make our concepts have, why should we fuss about what meanings our concepts currently do have?* One reason is that I like icing on my cake. But a second reason is that attributing meanings to our concepts often helps us to achieve benefits. It would run against the spirit of Pragmatic Conceptual Analysis to suggest that we abandon this beneficial practice. Instead we should ask what explication of CONCEPT-MEANING we should adopt to achieve the benefits of this practice more consistently. Doing this will help us to make useful progress not only in Philosophy of Mind, Philosophy of Language, and Cognitive Science, but also in our everyday attributions of concept-meaning.

Fortunately my desire for icing turns out to fit nicely with my commitment to follow the normative authority of Pragmatic Conceptual Analysis wherever it leads. For, in Stage 2 of our bootstrapping argument, we will use this normative authority to justify adopting a particular explication of our shared concept of concept-meaning.²⁴ Having explicated CONCEPT-MEANING in this way, we will then see, in Stage 3, that Pragmatic Conceptual Analysis is actually a semantically conservative tool for revealing concept-meaning, thus understood. In other words, Pragmatic Conceptual Analysis not only has normative authority, it also has *descriptive authority* – it reveals what we have epistemic and practical reason to think of as being the antecedent meanings of our concepts. And that's icing on our cake.

So we want to use Pragmatic Conceptual Analysis to explicate CONCEPT-MEANING. To do this we must first seek a 'job description' – an empirical analysis of how this shared concept has regularly delivered benefits. We've already seen empirical hypotheses about how reinforcement learning processes automatically and subconsciously focus our application of concepts onto practically relevant referents. Similar processes are probably employed by many non-human animals. But now we're asking about a uniquely human capacity (so far as we know): the capacity to direct a conceptual system at itself and

²⁴ One might worry whether we really have a concept of *concept*-meaning, distinct from our concept of *word*-meaning. I think we do. We can easily debate both what should count as a parent in contexts where we're considering hereditary illness, and what should count as a parent in contexts where we're considering choices in how a child will be reared. I construe these two debates as not being about the meaning of the word 'parent', but instead as being about the meanings of the two concepts which we regularly use this one word to express: something like BIOLOGICAL PARENT and CUSTODIAL PARENT. (If you want to resist this and insist that ordinary folk don't distinguish word-meaning from concept-meaning, then you are welcome to read my argument below as being about word-meaning – this would slightly restrict the scope of my conclusion, but wouldn't remove its significance.)

attribute meanings to particular concepts within itself. What regular benefits do we get from using a concept of concept-meaning in this way?

This is an empirical question and I will propose two empirical hypotheses in response. Everyday experience is sufficient to render these hypotheses quite plausible. However, they can, and eventually should, be subject to empirical testing.²⁵

First, our attributions of concept-meaning provide further help in the process of focusing our application of various concepts onto practically relevant referents.

H1. CONCEPT-MEANING has regularly delivered benefits by guiding our application of other concepts.

Second, given that all these mechanisms serve continually to focus our usage of our concepts in these ways, our attributions of concept-meaning allow us to capture robust and explanatorily useful patterns in the functioning of these mechanisms.

H2. CONCEPT-MEANING has also regularly delivered benefits by helping us to give useful explanations and predictions regarding people's use of other concepts.

We'll consider each hypothesis in turn.

2.2. Guiding Concept Application

Let us begin our pragmatic analysis of CONCEPT-MEANING by attending to (H1) the ways in which this concept regularly helps to guide our application of other concepts. I will argue that, to best continue this work, we should adopt an explication of CONCEPT-MEANING which would have us attribute meanings to various concepts in the same way that Pragmatic Conceptual Analysis does. This result shouldn't be surprising, for we've already seen that Pragmatic Conceptual Analysis has practical normative authority – i.e., that it serves as a good guide for our future application of various concepts.

Let's consider three particular ways that our concept CONCEPT-MEANING might help to guide our application of other concepts.

H1a. We sometimes use our knowledge of the meaning of a target concept to determine whether or not to apply that concept in a given circumstance. E.g., many people identify *rhombuses* by explicitly employing their knowledge that RHOMBUS means quadrilateral with four equal sides.

However, we typically *don't* use our concept CONCEPT-MEANING as we apply frequently-used concepts – instead we usually use recognitional skills that are more direct.

H1b. CONCEPT-MEANING helps us to develop direct recognitional skills that we use in applying our target concepts. E.g., after using one's knowledge of RHOMBUS' meaning to determine whether several things are rhombi, one quickly

²⁵ I hope that 'experimental philosophers' will someday move away from pumping naïve subjects' intuitions about fanciful cases and instead take up questions about how our concepts actually regularly deliver benefits. I outline such a 'pragmatic experimental philosophy' research program in (###).

develops the ability to recognize rhombi on sight without any direct appeal to one's knowledge of this concept's meaning.

This is naturally related to a third way in which this guidance occurs.

H1c. We learn socially how to apply many of our shared concepts. In teaching someone how to apply a shared concept correctly, we often draw upon knowledge of what that concept means. In some cases we state a concept's meaning explicitly. In other cases we give a watered-down version of the meaning ('a yak is a large cow-like animal') or we offer recognitional procedures that are consistent – or close enough to consistent – with what we take the concept to mean.

It is an empirical question whether we really get regular benefits in these three ways but reflection upon everyday experience strongly suggests that we do. (If you disagree, I encourage you to set up experiments to decide the matter.)

In each of these hypotheses, which meaning is attributed to a target concept determines how that concept will then be applied. When our concept CONCEPT-MEANING delivers benefits in these cases, this is typically by way of getting us to apply the target concept in beneficial ways. While this guidance might occasionally yield benefits in some novel way, much more often it will lead us to continue using target concepts to deliver benefits in the same ways that earlier uses of these concepts have regularly delivered benefits.

We now have a plausible 'job description' of the work our concept CONCEPT-MEANING has been doing for us: it has guided our application of other concepts so that they might deliver benefits in tried-and-true ways. Which explication of CONCEPT-MEANING would best allow us to continue this useful work? We need an explication that will attribute to each concept a meaning that, if employed as application conditions for that concept, would well sustain the regular ways in which that concept has been beneficially used. This characterization should sound familiar, for it is precisely the sort of meaning that Pragmatic Conceptual Analysis attributes to various concepts. Let us call this notion of meaning *pragmatic meaning*. (Note that the pragmatic meaning of any concept is the meaning that Pragmatic Conceptual Analysis would attribute to that concept. The pragmatic meaning of the Huron concept was scurvy; the pragmatic meaning of FREE ACTION is some sort of compatibilist understanding of free action; and the pragmatic meaning of CONCEPT-MEANING is pragmatic meaning itself.)

We have seen a job CONCEPT-MEANING regularly does for us and we have seen that *pragmatic meaning* is the formal notion best capable of doing this job. Hence Pragmatic Conceptual Analysis tells us to adopt *pragmatic meaning* as our explication of CONCEPT-MEANING. By the normative authority of Pragmatic Conceptual Analysis established in Stage 1, we therefore have practical and epistemic reason to adopt this explication.

This argument turned upon the fact that we use our concept CONCEPT-MEANING to further disseminate and exploit regular ways in which target concepts have already been delivering benefits. This usage is best sustained by explicating CONCEPT-MEANING in a way that leads us to think about our own concepts much as we naturally thought about the Huron concept of scurvy. Like the Huron concept, our own concepts are imperfect but useful tools for gathering information about important referents in our world. In attributing

meanings to these concepts, we do well to focus upon sustaining the regular benefits these concepts have delivered and not upon our sundry intuitions surrounding these concepts nor upon the all the ways we've tried applying them that haven't worked.

2.3. Explanatory and Predictive Uses

I have concentrated upon one use of our concept CONCEPT-MEANING, namely guiding us to apply other concepts in beneficial ways. I argued that this usage would be best sustained by adopting the notion of pragmatic meaning as our explication of CONCEPT-MEANING. One might wonder whether we would reach the same conclusion if we consider other uses of CONCEPT-MEANING. In particular one might want to focus on the second hypothesized usage mentioned above:

H2. CONCEPT-MEANING has also regularly delivered benefits by helping us to give useful explanations and predictions regarding people's use of other concepts.

There are close links between (H1) the task of guiding our usage of various concepts and (H2) the task of predicting how beneficial various attempts to use those concepts will be. Good guides must be good predictors and good predictors make good guides. We've already seen that Pragmatic Conceptual Analysis is a good guide: when we apply our concepts in accordance with the pragmatic meanings Pragmatic Conceptual Analysis attributes to them, we do optimally well at continuing patterns of beneficial concept usage. Since Pragmatic Conceptual Analysis is a good guide, it is also a good predictor – predicting that, so long as things continue normally, applying a concept in accordance with its pragmatic meaning will tend to be more beneficial than applying it in other ways. Given this close linkage, we should expect that the notion of pragmatic meaning will be well-suited not just to the guidance tasks described in H1, but also to the explanatory and predictive tasks described in H2.

Here's a second reason to expect this result. We saw in section 1.2 that mechanisms of positive reinforcement learning play a significant role in calibrating our use of concepts. These mechanisms try various ways of doing things, abandon those that don't work, and keep those that are regularly beneficial. In attributing a *pragmatic meaning* to a concept, we effectively highlight the ideal that these trial-and-error processes have been calibrating our use of that concept to approximate. For example, in saying the Huron concept meant scurvy, we highlight the fact that it was *scurvy* that was providing the reward signal in the Huron's learning process – it was *scurvy* that the Huron were continually learning to better recognize and better cure. Highlighting such an ideal is predictively useful because positive reinforcement learning mechanisms will ensure that concept-users will remain fairly close to this ideal. It is also explanatorily useful because it enables us to understand various cases of practical success as conforming to this ideal and various cases of practical failure as departures from it. E.g., when the Huron misapply their scurvy concept to, say, a case of muscular dystrophy, they will behave in a way that would be a great response to scurvy – a treatment rich in vitamin C – but this behavior likely won't be beneficial in response to other ailments.

The conclusions of this section match those of Section 2.2. We've considered explanatory and predictive tasks that we regularly call upon our concept of concept-meaning to do, and we've seen that, as was the case with the application-guiding tasks considered before, we

would need to adopt *pragmatic meaning* as our explication of CONCEPT-MEANING in order best to do these tasks. Hence, by the normative authority of Pragmatic Conceptual Analysis, we have practical and epistemic reason to adopt this explication.

Stage 3. Descriptive Authority

3.1. A Qualified Answer

In Stage 1 of the bootstrapping argument we established that Pragmatic Conceptual Analysis has normative authority: we have both practical and epistemic reason to adopt the explications it delivers.

In Stage 2 this normative authority licensed using Pragmatic Conceptual Analysis to explicate CONCEPT-MEANING. Pragmatic Conceptual Analysis delivered the notion of *pragmatic meaning* as its explication of this concept. Hence we have practical and epistemic reason to adopt the notion of *pragmatic meaning* as our explication of CONCEPT-MEANING.

Finally, in Stage 3, we return to the question with which we began: does Pragmatic Conceptual Analysis correctly reveal concept-meaning? We know that Pragmatic Conceptual Analysis reveals the pragmatic meanings of various concepts, and we've seen that we have practical and epistemic reason to adopt the notion of pragmatic meaning as our explication of CONCEPT-MEANING. Hence, the answer to our question is a qualified 'yes': Pragmatic Conceptual Analysis does indeed reveal concept-meaning, at least if you understand 'concept-meaning' in the way that we have practical and epistemic reason to understand it, namely as pragmatic meaning.²⁶ So, in addition to its practical and epistemic normative authority, Pragmatic Conceptual Analysis also has *descriptive authority* – it reveals what we have both practical and epistemic reason to think of as having been the meanings of our concepts all along.

3.2. Dispensing with the Qualifier

My bootstrapping argument gave a *qualified* answer to our initial question: yes, Pragmatic Conceptual Analysis does reveal concept-meaning, *if you understand concept-meaning in the way we have practical and epistemic reason to understand it*. One might wonder how much weasel-work this italicized qualification is doing. From the perspective of someone who hasn't yet come to think of CONCEPT-MEANING as meaning pragmatic meaning, my proposal might seem to demand a radical change in our shared conceptual framework – it demands that we start thinking that CONCEPT-MEANING means pragmatic meaning rather than whatever we used to think it meant. The fact that we have practical and epistemic reasons to make this revision may do little to mitigate the suspicion that I'm sneakily changing the subject, from a question about whether Pragmatic Conceptual Analysis preserves concept-meaning in our pre-explication sense of concept-meaning, to a question about whether it preserves concept-meaning in some new sense.

²⁶ People sometimes ask why they can't just as easily offer bootstrapping arguments for their own theories of concept-meaning. E.g., why can't someone argue that the descriptive-meaning of CONCEPT-MEANING is descriptive-meaning, so we should embrace descriptive-meaning as our explication of CONCEPT-MEANING? My response: even if some other theories of concept-meaning are self-ratifying in the sense needed to sustain Stages 2 and 3 of a bootstrapping argument, they still won't have the (Stage 1) normative authority that would make their ratifying themselves be a reason for us to ratify them too.

In a moment, I'll explain why I don't think my proposed explication actually would change the subject, but first let's consider how bad it would be if it did. If my proposed explication would change the subject then it must be that, currently – before we adopt my proposed explication – CONCEPT-MEANING means something other than pragmatic meaning. (E.g., maybe it means intuitive meaning.) We've seen that the notion of pragmatic meaning is optimally suited to the sorts of work we regularly call upon CONCEPT-MEANING to do. So whatever the current meaning of CONCEPT-MEANING is, if it isn't pragmatic meaning, it must be comparatively poorly suited to doing this work. Hence we have practical reason to lay aside this old meaning and to adopt pragmatic meaning as our new explication of CONCEPT-MEANING.²⁷ Furthermore, this change is not just *practically* advisable; it's what our tried-and-true strategies for achieving *epistemic* success call upon us to do. Hence if what I've offered is a change in subject, it should be a welcome change.

Typically, when we stipulate a change in a concept's meaning, we must proceed with great caution, for information we'd previously associated with that concept may not fit its new referent. Less caution is required when the proposed explication is the concept's pragmatic meaning. For pragmatic meaning is defined in such a way that applying a concept in accordance with its pragmatic meaning will preserve that concept's old patterns of regular beneficial usage. We can confidently continue using old information we had associated with the concept, for, even after the proposed explication, this information will fit the concept's referent well enough to continue bringing about benefits in the same old ways. So, if it's revisionary to adopt a concept's pragmatic meaning as its explication, at least we needn't worry about treating this as a revision.

This is reflected in our everyday epistemic practices. In section 1.2, I argued that common mechanisms of positive reinforcement learning are an effective implementation of Pragmatic Conceptual Analysis. We do positive reinforcement learning all day, every day, without worrying that it might bring about semantic revisions. Since our tried-and-true epistemic practices don't worry that Pragmatic Conceptual Analysis might lure us into making semantic revisions when we follow it subconsciously and automatically, we probably shouldn't worry that it'll lead us to make semantic revisions when we follow it self-consciously and deliberately.

We've seen that, if applications of Pragmatic Conceptual Analysis lead us to change the subject, this should be a welcome change – so welcome that we have practical and epistemic reason to make this change and not even treat it as a change. Given all this, one must wonder why anyone would want to insist that our concepts currently mean something other than the pragmatic meanings that our practical and epistemic reasons will lead us, without any ado, to say they mean.

²⁷ I should stress that I'm not asking any theorists entirely to abandon their favorite meaning-like notions – I'm just proposing that they stop using these notions as their explication of CONCEPT-MEANING. These theorists are welcome to introduce new theoretical concepts like CAUSAL-CONCEPT-MEANING or INTUITIVE-CONCEPT-MEANING or WHATEVER-SORT-OF-CONCEPT-MEANING-THEY-LIKE, and to use these new concepts to articulate the sorts of claims they thought they could articulate using our ordinary concept CONCEPT-MEANING. I have argued, however, that these other theoretical concepts will be relatively poorly suited to the sorts of work that we have called upon CONCEPT-MEANING to do, and I suspect they won't be well-suited to much other useful work either. We may introduce these other meaning-like notions into the marketplace of ideas and let them fare as they will. I suspect they'll end up relegated to the graveyard of well-defined but not-very-useful theoretical concepts, which is where I think they belong.

Still, staunch defenders of non-pragmatic views of concept-meaning might insist that Pragmatic Conceptual Analysis and its subconscious equivalents really do force changes in concept-meaning, even if our tried-and-true epistemic practices don't flag them as such. Of course, there are proposed theoretical understandings of 'concept-meaning' upon which this claim is true – e.g., everyday explicative practices probably do subtly reshape our intuitions and thereby change the 'intuitive meanings' of our concepts without our noticing it. The fact that we don't notice it, though, is evidence that this isn't actually what we mean by 'concept-meaning'. Instead this evidence weighs in favor of the hypothesis that whatever we're thinking of as 'concept-meaning' is something that typically isn't changed by our everyday explications – e.g., something like *pragmatic meaning*.

This argument might not be decisive, as the advocate of a non-pragmatic view of meaning can surely cook up some story to explain why we fail to notice all these supposed changes in meaning. However, such an advocate faces the heavy burdens of concocting this story in a non-*ad hoc* way and of offering positive motivation for her view even while we have practical and epistemic reason to shift our attention nonchalantly from her favorite notion of concept-meaning to pragmatic meaning.

Until given such a story, we should follow our practical and epistemic reasons and embrace pragmatic meaning as our explication of CONCEPT-MEANING. Having explicated CONCEPT-MEANING in this way, we see that, of course, CONCEPT-MEANING meant pragmatic meaning all along: even before we realized how this concept worked, it was already our imperfect, but still useful enough, tool for tracking pragmatic meaning.

It's perfectly appropriate for us to adopt this explication without acting as though we might have made a stipulative revision. We've already seen that we have practical and epistemic reason not to be especially concerned that this explication might have been revisionary. Furthermore, for most *social* purposes, we needn't treat this as a stipulative revision. The various members of our speech community have diverse ways of recognizing when our shared concepts apply²⁸ and we each continually use various sorts of learning to update our own recognitional capacities. Given all this diversity and change, it's usually pointless to try to track who has done which explications to their own tokens of our shared concepts. Instead, we each follow our practical and epistemic reasons in updating our own recognitional procedures on the fly, and we count upon the fact that others will have reason to end up explicating their own tokens of our shared concepts in the same ways.

So, suppose other people ask us the question with which we began: "Does Pragmatic Conceptual Analysis correctly reveal concept-meaning?" We ordinarily can't be bothered by worrying about whether they might have done a bit less explication on their own tokens of our shared concept of concept-meaning than we have on ours. Instead, it is ordinarily felicitous for us to answer, "Yes", and count on them to do the normatively appropriate explication themselves (or ask us for confirmation) before they go too far astray using our new information with their old explication. We have perfectly good error-correction procedures that can kick in if their explications fail to keep up with ours, and we can use those procedures in the special circumstances where they're needed. But for ordinary circumstances it is perfectly appropriate for us to dispense with the qualifier and say that Pragmatic Conceptual Analysis reveals the meanings of our concepts – full stop.

²⁸ This point is emphasized by Millikan (1998, 2000).

References.

- Appiah, Anthony. 1986. "Truth Conditions: A Causal Theory." In *Language, Mind and Logic*, Thyssen Seminar Volume, Jeremy Butterfield (ed.) Cambridge University Press, pp. 25-45.
- Blackburn, Simon. 2005. "Success Semantics". In *Ramsey's Legacy*, H. Lillehammer and D. H. Mellor, eds. Oxford University Press.
- Boyd, Richard. 1988. "How to be a Moral Realist," in *Essays on Moral Realism*, ed. Sayre McCord. Cambridge University Press, pp. 181-228.
- Burge, Tyler. 1979. "Individualism and the Mental." in P. French, T. Uehling, and H. Wettstein, eds., *Studies in Metaphysics*. University of Minnesota Press.
- Carnap, Rudolf. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Chalmers, David. 2002. "The Components of Content" in David J. Chalmers ed., *Philosophy of Mind: Classical and Contemporary Readings*. New York: Oxford UP, pp. 608-633.
- _____. (2004) "The Foundations of Two-Dimensional Semantics" in M. Garcia-Caprintero and J. Macia, eds. *Two-Dimensional Semantics: Foundations and Applications*. Oxford University Press.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Oxford UP.
- Davidson, Donald. 1984. *Inquiries into Truth and Interpretation*, Oxford: Clarendon.
- Dennett, Daniel. 1987. "Making Sense of Ourselves" in his *The Intentional Stance*. MIT Press.
- _____. 1991. *Consciousness Explained*. Boston: Little, Brown.
- Dretske, Fred. 1988. *Explaining Behavior*. MIT Press.
- Erev, Ido. 1998. "Signal Detection by Human Observers: A Cutoff Reinforcement Learning Model of Categorization Decisions Under Uncertainty." *Psychological Review*. 105, 280-98.
- Evans, Gareth. 1973. "The Causal Theory of Names." *The Philosophy of Language*. Ed. A.P.Martinich. New York: Oxford UP, 1996.
- Fodor, Jerry. 1990. *A Theory of Content and Other Essays*. MIT/Bradford.
- Frege, Gottlob. 1892. "On Sense and Nominatum" in A. P. Martinich, ed., *The Philosophy of Language*. 3rd ed. Oxford: Oxford University Press, 1996. pp. 186-98.

- Griffiths, Paul. 1997. *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.
- Henrich, J. and R. McElreath. 2007. "Dual Inheritance Theory: The Evolution of Human Cultural Capacities and Cultural Evolution." In *Oxford Handbook of Evolutionary Psychology*, edited by Robin Dunbar and Louise Barrett. Oxford UP.
- Haslanger, Sally. 1999. "What Knowledge Is and What It Ought To Be: Feminist Values and Normative Epistemology." in *Philosophical Perspectives* 13: 459-480.
- Henderson, David. (1994) "The Principle of Charity and the Problem of Irrationality (Translation and the Problem of Irrationality)" in *Readings in the Philosophy of Social Science* ed. Michael Martin and Lee C. McIntyre. Cambridge, Mass.: The MIT Press. 323-341.
- Hilbert, David. 1987. *Color and Color Perception*, Stanford, Calif.: C.S.L.I.
- Hull, David. 1978. "A Matter of Individuality." *Philosophy of Science* 45, 335-60.
- James, William. 1906. "What Pragmatism Means."
<http://www.marxists.org/reference/subject/philosophy/works/us/james.htm>
- Johnston, Harry. 1912. *Pioneers in Canada*.
<http://www.blackmask.com/thatway/books162c/pican.htm>
- Joyce, James. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65, 575-603.
- Kornblith, Hilary. 2002. *Knowledge and Its Place in Nature*. Oxford UP.
- Kripke, Saul. 1972. *Naming and Necessity*. Cambridge: Harvard UP.
- Lewis, David. 1983. "Radical Translation" in his *Philosophical Papers*. New York: Oxford University Press. 108-118.
- _____. 1984. "Putnam's Paradox." *Australasian Journal of Philosophy* 62: 221-236.
- Margolis, Eric and Stephen Laurence (eds). 1999. *Concepts: Core Readings*. Cambridge, MA: MIT Press.
- Melville, Herman. 1851. *Moby Dick*.
- Millikan, Ruth. 1984. *Language, Thought, and Other Biological Categories*. MIT Press.
- _____. 1998. "A common structure for concepts of individuals, stuffs, and basic kinds: More mama, more milk and more mouse." *Behavioral and Brain Sciences* 22.1, 55-65.
- _____. 2000. *On Clear and Confused Ideas*. Cambridge, UP.

- Papineau, David. 1987. *Reality and Representation*. Blackwell.
- Pascal, Blaise. 1910. *Pascal's Pensées*, translated by W. F. Trotter.
- Pollock, John and Joseph Cruz. 1999. *Contemporary Theories of Knowledge*. 2nd edition, Lanham, MD: Rowman & Littlefield.
- Putnam, Hilary. 1973. "Meaning and Reference." *The Philosophy of Language*. Ed.A.P.Martinich. New York: Oxford UP, 1996.
- Quine, W.V.O. 1960. *Word and Object*. Cambridge: The MIT Press.
- _____. 1969. "Epistemology Naturalized," in his *Ontological Relativity and Other Essays*. New York, Columbia University Press.
- Ramsey, Frank. 1927. "Facts and Propositions." In *The Foundations of Mathematics, and Other Logical Essays*. R.B. Braithwaite, ed. London: Routledge and Kegan Paul, 1931, pp. 138-55.
- Rupert, Robert. 1999. "The Best Test Theory of Extension: First Principle(s)". *Mind & Language* 14: 321-55
- Strawson, P.F. 1950. "On Referring." in A. P. Martinich, ed., *The Philosophy of Language*. 3rd ed. Oxford: Oxford University Press, 1996. pp. 215-30.
- Thorndike, Edward. 1898. "Animal intelligence: An experimental study of the associative processes in animals." *Psychological Review Monograph Supplement* 2, 1-109.
- _____. 1911. *Animal intelligence*. New York: Macmillan.
- Woodward, James. 2004. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.
- Whyte, Jamie. 1990. "Success semantics." *Analysis* 50: 149–57.