

Toward a Theory of Concept Mastery*

Gabriel Oak Rabin

Draft of October 2015

Do not cite or distribute without permission.

Abstract: This paper investigates the question “Under what conditions does a thinker fully understand, or have mastery of, a concept?” A thinker can possess, and think thoughts with, concepts he or she does not fully understand. I argue against three views of concept mastery, according to which mastery is a matter of holding certain beliefs, being disposed to make certain inferences, or having certain intuitions. None of these attitudes is either necessary or sufficient for mastery. I propose and respond to objections to my own “recognition view” of the conditions under which a thinker has mastery of a concept.

1 Introduction

Concepts are the basic units of thought. The thought OSTRICHES LIKE CHOCOLATE is composed of the concepts OSTRICHES, LIKE, and CHOCOLATE. An agent *possesses* a concept when he or she can think thoughts of which

*Thanks to Denis Buehler, Tyler Burge, David Chalmers, Mark Greenberg, Tamar Weber, and audiences at the American University in Beirut, Arch Research Centre at the University of St Andrews, The Australian National University, the University of Edinburgh, the University of Melbourne, the University of Sydney, the 2014 Meeting of the Australasian Association of Philosophy, and the UCLA Albritton Society for helpful comments and criticism.

the concept is a component. An agent has *mastery* of a concept when he or she fully understands that concept. One can possess a concept without having mastery. Many people use technical concepts that have worked their way into public consciousness without fully understanding them (examples include DARK MATTER, CHAOS THEORY, CLOUD COMPUTING, TRANSISTOR, ID, FRACKING, and HEDGE FUND). Any theory of concept mastery must answer the following fundamental question: “Under what conditions does an agent have concept mastery?” This is the question I hope to answer, at least in part, in this paper.

I will argue against three versions of *the endorsement view*, according to which mastery is a matter of endorsing certain propositions or inferences. The endorsement view comes in at least three forms: “the belief view”, “the inference view”, and “the intuition view”, according to which endorsement involves believing, inferring, or intuiting in accordance with certain propositions or inferences. I will argue for a different approach to concept mastery, *the recognition view*, according to which mastery is a matter of recognizing certain rules as governing the use of the concept. Recognition is compatible with a failure to endorse.

I now offer a roadmap of the paper. In section 2 (“Preliminaries”) I clarify terminology, lay out assumptions, argue briefly against skepticism about concept mastery, and canvas four potential views about what it is to master a concept. Each of the three views against which I argue maintains that some set of attitudes (beliefs, inferences, or intuitions) is necessary and sufficient for mastery of a concept. In sections 3 (“Deviant Masters”) and 4 (“The Special Status Problem”) I argue against the necessity and sufficiency

(respectively) of these attitudes for concept mastery. The failures of these views point the way toward a successful theory, which I present in section 5 (“The Recognition View”). Section 6 (“Four Objections”) considers and replies to objections to the recognition view.

2 Preliminaries

2.1 Terminology and assumptions

In this section, I’ll lay out assumptions, clarify terminology, and sketch a framework in which to commence inquiry. First, I use the term ‘proposition’ to describe a structured composite of concepts. For example, the proposition OSTRICHES LIKE CHOCOLATE is composed of the concepts OSTRICHES, LIKE, and CHOCOLATE. I use small capitals, e.g. OSTRICH, to express elements at the conceptual level, either concepts or complex propositions. Following Frege [1952/1892], I also use the term ‘thought’ to denote propositions so described. I adopt a broadly Fregean approach to concepts and propositions. This approach is not intended to be controversial. Inevitably, due to prior commitments about the nature of propositions, some readers will take it as such. Readers who bridle at my use of ‘proposition’ should substitute the word ‘propersition’ for every occurrence of ‘proposition’. Propersitions are the complete thoughts I am currently describing. I assume that one can, and for many purposes should, type mental states using propositions (propersitions) so described.

A certain class of mental states (the so-called “propositional attitudes”) can usefully be categorized by a combination of a *mode* and a proposition. The mode characterizes (at least one of) the thinker’s attitude(s) toward the

proposition. Examples of modes include belief, desire, and fear. Examples of complete propositional attitudes include a belief that ostriches like chocolate, a desire that ostriches like chocolate, and a fear that ostriches like chocolate. I don't deny that there are many other ways to type and describe mental states and propositions, or that those ways are useful, informative, or illuminating. I assume only that the locutions I've described constitute one useful way of thinking and talking about contentful mental states. Something like the above approach should be compatible with and fruitful for a wide variety of philosophical and psychological projects. I'll assume this approach throughout the paper.

I assume that there is a distinction between merely possessing a concept (i.e. being able to take some attitude toward a proposition that contains the concept) and fully understanding that concept.¹ Not uncontroversially, I assume also that, sometimes, thinkers entertain thoughts containing concepts they do not fully understand, and thus possess those concepts without mastery of them. Theorists of content and of concepts tend to assume, often implicitly, without argument, and without recognizing that any such assumption has been made, that possession of a concept (i.e. the ability to think thoughts containing the concept) entails mastery of that concept. Often they fail to note any distinction at all between possession and mastery. The difficulty is partly due to ambiguity in the terminology of "having/possessing a concept". That locution can be used to express either a relation of possession or a relation of mastery between a thinker and a concept. Here, I reserve the locution "x possesses the concept C" to mean that x can think thoughts that

¹For further discussion of this distinction, cf. Bealer [2002a]: 221-222, Burge [1979, 1986], Greenberg [MS]: 13-4, Greenberg [2001]: 28-9, Greenberg [2014], .

contain C. x may or not may have mastery of C.

In so far as theorists do recognize the possibility of concept possession without mastery, they often chalk up the phenomenon to and add a proviso about “deference”, and then continue to ignore the possibility of possession without mastery. Here, I do not wish to tackle issues about the relation between concept possession and concept mastery and their relations to the theory of content and of concepts more generally. (I point the reader to the discussions of Greenberg [MS, 2001, 2014], Burge [1979, 1986], and Peacocke [1992, 1999, 2008]). Instead, I will simply assume that thinkers can possess concepts that they have not mastered. Given that assumption, I then seek to answer the question, “What does it take for a thinker to master, or fully understand, his or her concept?”

I take no official position on what is required to possess a concept. I require only that whatever is required to possess a concept, it does not always entail that one have mastered the concept. However, I believe (but will not assume) that the requirements on concept possession are relatively minimal. It’s easy to acquire a concept, especially when that concept is expressed by a public language term. Hearing someone utter a term that expresses a concept not already possessed, and then using that term and concept in thought, without intending to appropriate the term and give it new meaning, will often be sufficient for acquisition of the concept. For example, if I hear a doctor declare that the patient has lupus, and begin believing that the patient has lupus, then I can come to possess, but not master, the concept LUPUS. This is so despite my ignorance regarding LUPUS. I know little about LUPUS other than that it refers to some type of ailment.

Most controversially, I assume a distinction between “core” elements associated with a concept and “non-core” elements. These elements might be propositions, inferences, mental transitions, rules, or something else. Core elements are in some sense especially important to the meaning of the concept. They play a special role in the concept’s having the representational qualities that it does and being the concept that it is. One might further claim that core elements partly or wholly constitute the nature of the concept. For example, I take the proposition ALL BACHELORS ARE MALE to be a “core” proposition for the concept BACHELOR. ACHILLES IS A BACHELOR is non-core. One might also maintain that core propositions are “analytic” or contribute to *a priori* justification. I remain agnostic on this count. To accept the distinction between the core and the non-core, one must resist, at least to some degree, the arguments of Quine [1951] against the analytic/synthetic distinction.

2.2 Cores and Conceptual Roles

According to a broad philosophical tradition called the “conceptual role approach” to mental content and concepts, a concept’s role in thought, and especially its connections to other concepts, plays a vital role with respect to the concept. For example, according to Greenberg [MS]’s characterization, conceptual role theories (which he calls ‘inferential-role theories’) maintain that “what makes it the case that a thought involves a particular concept is that the thinker deploys a disposition to move between mental states that is appropriately isomorphic to the concepts canonical or defining connections to other concepts.” (21) The main rival to conceptual role theories are informational/covariational accounts, which maintain that a concept’s reference,

rather than its relations to other concepts, plays the vital role.

My use of the phrase ‘a/the vital role’ is deliberately vague, because extant theories tend to run together a variety of theoretical roles that conceptual or inferential roles, or transitions between mental states more broadly, might play. I envision at least six. First, one might claim that conceptual roles partly or wholly individuate concepts. Second, one might claim that conceptual roles partly or wholly constitute the nature of the concept. Third, one might claim that a concept’s role determines the extension of a concept. For example, a theory might say that the extension of the concept is that which makes the concept’s canonical transitions truth-preserving. Fourth, one might tie concept possession to conceptual role, in the sense that possession of a concept requires having a mental item that plays the appropriate role. Fifth, one might further claim that conceptual roles make it the case that a particular mental state or thought contains one concept rather than another. (These last two are closely connected, but there is a difference. The former concerns a general ability (possession), while the latter pertains to the content of a particular mental state.) Sixth, one might claim that conceptual roles determine mastery of the concept, in the sense that one has mastery of the concept if and only if one appropriately grasps the concept’s conceptual role or canonical transitions.

My project and sympathies fall with the conceptual role tradition, broadly conceived. My acceptance of a distinction between “core” and “non-core” elements, which I described as “propositions, inferences, mental transitions, rules, or something else”, places me in this camp. The core elements are something like the “canonical transitions” on which conceptual role theories

rely. All the accounts of concept mastery that I will discuss fall broadly within the conceptual role approach.

It's worth laying my views on the table, partly to help the reader understand the angle from which I approach this inquiry. I believe that canonical transitions partly, but not wholly, individuate concepts (reference does some individuation as well). I believe that canonical transitions partly constitute the nature of the concept. I do not believe that canonical transitions determine extension. I tie neither concept possession nor what makes a thought contain one concept rather than another, to conceptual role. My favored account of concept mastery, the "recognition view" (section 5), says (roughly) that mastery of a concept amounts to understanding of the canonical transitions. I will not assume any of these beliefs in the discussions to follow. I will argue for (not assume) the claim that concept mastery involves understanding of canonical transitions. The three elements of conceptual role approaches that I endorse (concept individuation, nature and mastery) fit together. From the claim that concepts are partly or wholly individuated in terms of canonical transitions, it's a short step to the claim that those transitions form part of the concept's nature. It's then natural to claim that understanding the concept is a matter of understanding that nature, which entails understanding those canonical transitions.²

2.3 Precising the Notion of concept mastery

In this section, I clarify the target notion: mastery, or full understanding, of a concept. A thinker has mastery of the concept FLYING BUTTRESS ex-

²Greenberg [2014]'s account of concept mastery manifests a similar connection between nature and understanding of that nature. He writes (1,7) that mastery of a concept is a matter of "having a practically available understanding of its nature or essence."

actly when she fully understands FLYING BUTTRESS. I believe that the notion of full understanding, or having mastery, of a concept, is relatively pre-theoretical and everyday. We express the relation using locutions like ‘she fully understands the concept flying buttress’, ‘she grasps the meaning of the concept flying buttress’, and ‘she fully understands the meaning of ‘flying buttress’’.³ Mastering a concept is not an impossible feat, or an idealized state achievable only by hypothetical thinkers. You and I have mastery of many concepts.

Mastering a concept is not merely a matter of knowing many truths involving the concept. One might know many truths containing CHAIR, but if one fails to realize that chairs are for sitting, then one fails to master CHAIR. Mastery is tied to grasping certain important or key insights about the concept. Mastery of a concept is compatible with ignorance of important truths about that to which the concept refers. For example, I believe that mastery of the concept WATER is compatible with ignorance of the fact that water is

³We must be careful here. It’s possible that fully understanding the linguistic meaning of a word and fully understanding the concept it expresses could come apart. Burge [1990] claims that this is possible, and attributes to Gottlob Frege a similar view. The example Burge offers involves the mathematical concept LIMIT from the calculus. In the early days of the development of the calculus, both Isaac Newton and Gottfried Leibniz used the notion of a limit. Both possessed the concept LIMIT, but neither thinker had a clear grasp on it. The linguistic meaning of the word ‘limit’, however, is determined by how individual speakers have used the term. With respect to the linguistic usage of ‘limit’, Newton and Leibniz were the experts. They fully understood the linguistic meaning of ‘limit’. Yet they failed to fully understand the concept LIMIT that the word expressed. If this description of the case is accurate, then it seems possible to fully understand a term’s linguistic meaning yet fail to fully understand the associated concept. We might take the expression ‘x fully understands the meaning of term t’ to have two uses. It can be used to say that x fully grasps the linguistic meaning of t. Or it can be used to say that x fully grasps the concept that t expresses. Only when ‘x fully understands the meaning of t’ is used in the first sense does mastery of a concept come apart from full understanding of the meaning of a linguistic item. Throughout this paper, I’ll use the expression ‘fully understanding the meaning of term t’ exclusively in the second sense.

H₂O. Someone who knows that water is H₂O knows more about water than someone who does not. She also better understands water (the stuff) and its nature. But I deny that she better understands the concept WATER. That water is H₂O is a deep, necessary, and constitutive truth about water. But H₂O-ness is not part of the concept WATER itself.

Consider another case.⁴ Suppose that it turns out that all chairs are, in fact, artifacts left by an alien civilization. By an odd quirk of physics, the alien technology is necessary to create any object that can be sat on like a chair. Nothing can be a chair unless it is an alien artifact. If this is so, then “all chairs are alien artifacts” will be a deep and important truth about the nature of chairs. Nevertheless, I claim that one can fully understand and master the concept CHAIR without knowing about the chair-alien artifact link. Such masters of CHAIR will fail to fully understand chairs and their nature. But their understanding of the concept CHAIR need not be lacking. I take these claims about WATER and CHAIR to be intuitive. But even if one finds them controversial, they should help the reader grasp the notion of concept mastery as I conceive it.

2.4 Against Skepticism

In developing a theory of concept mastery, I take as my starting point the seminal discussions of Burge [1979, 1982, 1986]. Burge deserves much credit for bringing into the philosophical arena the phenomenon of *incomplete understanding*. Consider Burge [1979]’s example of Alf. Alf approaches his doctor, complaining of arthritis in his thigh. Alf’s doctor corrects him, pointing

⁴This case is based on a thought experiment in Burge [1986]: 263-64. I have altered the case and put it to a different use.

out that one cannot have arthritis in the thigh. Arthritis, by definition, affects only joints. Alf accepts the correction, admits that he was mistaken to believe that he had arthritis in his thigh, and modifies his usage of ‘arthritis’ and ARTHRITIS accordingly. The important point here is that Alf was able to think the thought I HAVE ARTHRITIS IN MY THIGH, and thereby possessed the concept ARTHRITIS, while significantly misunderstanding the concept. Alf incompletely understood ARTHRITIS. Similar examples can be cooked up for almost any concept whatsoever.

If agents can possess concepts while incompletely understanding them, it’s natural to ask what agents must do to possess a concept and completely understand it - i.e. to have concept mastery. It’s natural to think that such complete understanding is attainable by everyday people. Not everyone is like Alf.

Some will be skeptical that there is any such thing as “mastery” or “full understanding”, or that such notions are theoretically significant or useful.⁵ I won’t tackle such skepticism head on, but I want to make a few remarks to clarify what full-blooded skepticism about concept mastery must be and to lay out some considerations against it.

One can understand better or worse. Understanding is, in general, a matter of degree. Thus mastery of a concept will be a matter of degree. We might rank hypothetical understanders of a concept according to their understanding. On one end of the spectrum we have agents who meet only the minimal requirements (whatever they may be) for possession of the concept. On the other end of the spectrum we have agents who are omniscient with

⁵Among the skeptics I count Ball [2013] and Williamson [2008, draft], [p.c.].

respect to the concept. They know every true proposition containing the concept and know all the important truths about the thing in the world to which the concept refers.

The strongest form of skepticism about concept mastery maintains that there are no theoretically significant breaks, marks, or groupings among the hypothetical agents so ranked. Instead there's a continuous spectrum. A weaker form of skepticism admits that there are theoretically significant groupings, but none of those groupings corresponds to the notion of concept mastery as I have described it. One initially tempting form of skepticism about concept mastery goes something like this. Every concept occurs in a wide variety of truths. Some of these truths are known by almost everyone who uses the concept. They're common knowledge. We use the term 'fully [completely] understands the concept C' as a label for those individuals who know all the common knowledge truths involving C. On this view, there is such a thing as concept mastery, but it marks only a linguistic accident.

I deny all three forms of skepticism. I believe that there are some theoretically significant breaks, marks, or groupings among the hypothetical agents described above. I believe that one (perhaps rough and ready) way to informatively group the agents is to divide them into those that have mastered the concept and those that have not.

I offer a brief argument against the skeptic. It relies on two claims. First claim: Incomplete understanding occurs. Second claim: With respect to a given concept, there is a distinction between core truths and non-core truths. "All bachelors are male" is core vis-a-vis BACHELOR. "No bachelors live on Mars" is not. Between two agents who are otherwise the same, there's

a big difference, with respect to understanding of BACHELOR, between (i) the agent who fails to accept that all bachelors are male while accepting that no bachelors live on Mars and (ii) the agent who fails to accept that no bachelors live on Mars but accepts that all bachelors are male. The second understands BACHELOR much better than the first. Once we have a theoretically significant distinction between core and non-core truths, we can define up a theoretically significant level of concept mastery according to which mastery involves having a “grip” on the core truths. (We might also use core inferences, rules, or what have you). This tells against both the strong and weak forms of skepticism above.⁶ The linguistic form of skepticism I described is defeated by the above argument in combination with the observation that both “all bachelors are male” and “no bachelors live on Mars” are common knowledge.

One might resist the above arguments by maintaining that most concepts are not like BACHELOR. According to this line, BACHELOR permits a distinction between core and non-core elements, but most, and perhaps almost all, concepts do not. I admit that BACHELOR is special in certain respects, and in some ways not representative of concepts more generally. BACHELOR has a particularly “strong” core. One can reasonably maintain that the core elements determine the reference of BACHELOR, provide necessary and sufficient conditions or bachelor-hood, or give the meaning of BACHELOR. For many (perhaps most) concepts, analogous claims will not be as plausible.

But, importantly, many concepts other than BACHELOR yield core elements. Consider some examples. CHAIR: chairs are for sitting. ARTWORK:

⁶Greenberg [2001]: 134, Greenberg [2014]: 7-11 implements a similar strategy in arguing against skepticism about concept mastery.

artworks are created. MUG: mugs function to hold liquid. WINE: wine is made from grapes. BLANKET: blankets have a large surface area to thickness ratio. CAVE: caves are enclosed spaces. SWORD: a sword has a blade. The method of examples cannot possibly show that even a significant minority of all concepts have cores. But hopefully the examples alleviate some of the fear that BACHELOR is idiosyncratic in having any core at all, and lend some support to the idea that a significant portion of concepts have a core. The reader can further satisfy him or her self by doing a few cases on his or her own. My discussion of concept mastery and full understanding will not apply to those concepts which do not admit of a core vs. non-core distinction. (I do not assume that there are any such concepts). From here, I will speak as if all concepts have a core vs. non-core distinction, but one can take my comments as restricted to those concepts which do.

In sum, if one (a) believes that incomplete understanding occurs, and (b) resists the arguments of Quine [1951] enough to believe in a distinction (even a rough distinction) between “core” and “non-core” elements of a concept, then one already has a theoretically significant notion of concept mastery (at least for whichever concepts admit of the core vs. non-core distinction). To resist motivation (a) is to reject the tremendously successful arguments of Burge [1979, 1982, 1986], the conclusions of which have since become philosophical orthodoxy. To resist motivation (b) is to deny ALL BACHELORS ARE MALE any special status, vis-a-vis BACHELOR, above and beyond JAVIER IS A BACHELOR. Few are willing to pay either price. If they do not, they deny the skeptical view.

2.5 Four Theories of Concept Mastery

I've already laid out the distinction between core and non-core elements relative to a concept. The elements might be propositions, inferences, rules, or something else. With the core/non-core distinction in place, a view naturally suggests itself. Fully understanding a concept, i.e. having mastery of that concept, entails having some type of "grip" on that core. This much leaves room for variation. Theories of concept mastery can vary according to what they take the grip to consist in. I consider four views.

1. the belief view
2. the inference view
3. the intuition view
4. the recognition view

I plan to argue against the first three views (belief, inference, and intuition), and in favor of my own recognition view. But first let's get the views on the table.

According to the belief view of concept mastery, an agent has mastery of concept *C* if and only if he/she believes all of *C*'s core propositions. On this view, the elements of the core are propositions. To grip the core elements is to believe them. According to the inference view of concept mastery, an agent has mastery of concept *C* if and only if he/she is disposed to infer in accordance with all of *C*'s core inferences. On this view, the elements of the core are inferences. To grip the core elements is to be disposed to infer in

accordance with them. Proponents of the inference view include Peacocke [1992, 1999].⁷

The belief and inference view are similar. Let’s consider an example of each. Suppose that “all bachelors are male” and “all bachelors are unmarried” are the core truths for BACHELOR. The belief view claims that an agent has mastery of BACHELOR if and only if she believes both that all bachelors are male and that all bachelors are unmarried. The inference view claims that there are two core inferences associated with BACHELOR. We might express them as follows:

$$\frac{x \text{ is a bachelor}}{x \text{ is male}} \qquad \frac{x \text{ is a bachelor}}{x \text{ is not married}}$$

The variable ‘x’ represents the fact that the agent is willing to infer from “x is a bachelor” to “x is male” for any x. The inference view claims that an agent has mastery of BACHELOR if and only if he or she is disposed to infer, for any x, from x IS A BACHELOR to x IS MALE and from x IS A BACHELOR to x IS NOT MARRIED.

According to the intuition view of concept mastery, an agent has mastery of concept C when he/she has intuitions that the core propositions are true (Bealer [2002a]: 221-230).⁸ On this view, the elements of the core are

⁷Peacocke [1992] offers a theory of what he calls “concept possession”. However, Peacocke does not mean what I mean by ‘concept possession’ or ‘possess(es) a concept’. Peacocke (29-30) states that a thinker can consider propositions, and entertain thoughts, that contain the concept while failing to possess the concept. On my use of ‘possess a concept’, to possess a concept just is to be able to entertain thoughts or propositions that contain the concept. Thus Peacocke does not mean what I do by ‘possess a concept’. His theory of “concept possession” is a theory of a more demanding relation: what I call ‘concept mastery’.

⁸Bealer calls mastery of a concept “determinate possession”. His theory is more nuanced, and more complicated, than the simplistic intuition view that I consider. I will not

propositions. To grip the core elements is to intuit their truth. Importantly, intuiting the truth of P is compatible with failing to believe that P, or with believing not-P. An agent might believe P on the basis of reliable testimony, while simultaneously feeling - i.e. intuiting - that P is true.

The intuition view relies on a relation of intuition between an agent and a proposition. There is a rich philosophical literature on the nature of intuition. I wish to be as non-committal as possible regarding those debates. Some maintain that intuition is its own sui generis mental attitude (Bealer [2002a,b]). Others claim that intuition is a form of belief (Sosa [1996, 1998], Williamson [2004, 2008]). Even if intuition is a form of belief, the intuition view will, most likely, not reduce to the belief view, because the belief view claims that concept mastery requires belief *simpliciter*, whereas the intuition view claims that concept mastery requires a particular form or type of belief (whichever type of belief intuitions are).

According to the recognition view of concept mastery, an agent has mastery of concept C if and only if he/she takes all the core elements to govern the use of C. The view is liberal regarding the nature of core elements. To grip a core element is to take it to govern the use of the concept. In section 5.2 (“Taking to Govern Use”) I explain what it is to take something to govern the use of the concept.

There are a variety of dimensions of variations even among these views. One could claim that mastery requires only a grip on some, perhaps a majority, of the elements in the core. Or one might hybridize views, claiming

explain his view in full here. However, his theory of concept mastery is a form of the intuition view, shares its same general shape, and has the same advantages and disadvantages that are relevant here.

that mastery is a matter of believing some propositions, inferring in accordance with others, and intuiting the truth of a third class. For simplicity's sake, I ignore these possibilities. I believe they do not significantly alter the dialectic.

3 Deviant Masters: Against the Necessity of Belief, Inference, or Intuition for concept mastery

3.1 Overview

In section 3 I argue against the necessity of belief in, inference in accord with, or intuiting the truth of, core elements for concept mastery. The argument revolves around the possibility of non-standard, or “deviant”, masters of the concepts in question. The deviant masters fully understand the concept but neglect to believe, infer, or intuit in accordance with core elements.

The inference view will be my representative stalking horse. In sections 3.2-3.5 I discuss only the inference view. The belief and intuition views can be defeated by arguments exactly analogous to those I discuss. The modifications required are minimal.

3.2 The Argument

In this section, I argue against the inference view. The basic problem with the view is that an agent can fully understand a concept (i.e. have concept mastery) without being disposed to infer in accordance with core inference patterns. Being disposed to infer in accordance with a concept's core inferences is not necessary for mastery of the concept.

Consider the concept *BOCHE*, a racial pejorative applied to Germans *circa*

World War I. BOCHE has two core inference rules.

$$\frac{x \text{ is German}}{x \text{ is a boche}} \qquad \frac{x \text{ is a boche}}{x \text{ is cruel}}$$

If a thinker is disposed to make both these inferences, then she is disposed to infer from “x is German” to “x is cruel”. According to the inference view, an agent who has mastered BOCHE must be disposed to infer from “x is German” to “x is cruel”.

Here’s the counterexample. Consider Agnes, who at one point in time fully understood BOCHE. She was disposed to infer, and often did infer, in the manner just described. However, during her travels in Germany, Agnes dined, lived, and fraternized with many Germans. She made German friends. As a result of her experiences, Agnes ceased to be disposed to infer from a subject’s German-hood to his or her cruelty. Agnes stopped believing that all Germans are cruel. However, Agnes still realizes that the inference patterns above play a central role in relation to the meaning of BOCHE and in the use of ‘boche’. This thought experiment yields the following argument against the inference view.

(P1) If the inference view is true, then Agnes has lost her mastery of BOCHE.

(P2) Agnes has not lost her mastery of BOCHE.

(C) Therefore: the inference view is false.

Premise (P1) follows from the statement of the inference view in combination with the description of the thought experiment. Premise (P2) is

tremendously plausible. Agnes understands BOCHE at least as well as, and perhaps better than, her unenlightened co-linguals.⁹ Agnes still recognizes the inference patterns as rules that govern the use of BOCHE. She simply refuses to obey them. The argument is valid. The conclusion follows: the inference view is false, because being disposed to infer in accordance with core inferences is not necessary for concept mastery.

3.3 Response 1: Agnes remains disposed to infer

Proponents of the inference view might respond by denying premise (P2), claiming that Agnes remains disposed to infer from “x is German” to “x is cruel”. A disposition need not always manifest itself; dispositions can be masked. A wine glass wrapped in protective layering remains fragile, and retains the disposition to break, even though it does not break when struck. The disposition to break is *masked* by the protective wrapping. The proposed response to the Agnes case claims that Agnes retains the disposition to infer in accordance with BOCHE’s core, but the disposition is masked, possibly by her belief that not all Germans are cruel.

Two points should be made against this response. First, the move is unmotivated and *ad hoc*. *Prima facie*, Agnes lacks the disposition. To maintain otherwise one must posit reasons in favor of her having a masked disposition. Proponents of the inference view have offered no such reason, other than that a masked disposition supports their favored theory. Second, a version of the Agnes case can be cooked up in which Agnes clearly lacks the disposition. Whatever the criteria are for having the disposition to infer, we can imagine a case in which Agnes fails to meet those criteria yet has explicit knowledge

⁹Williamson [2009]: 141 seems to agree.

of the inference rules governing the use of BOCHE. We can generate such a case as long as it's possible to know what the inference rules governing a concept are without being disposed to obey them. This is possible, so we'll be able to generate counterexamples to the inference view, whatever opponents claim about what it is to have a disposition to infer.

3.4 Response 2: boche is not a concept

BOCHE is an odd concept. Its inference rules are not truth-preserving. Someone who follows its constitutive inference rules will be led, quite rapidly, into error. (Few Germans are cruel). Such concepts can be called 'defective'. Proponents of inference-style views have sometimes claimed that there are no concepts like BOCHE (e.g. Horwich [2010]: 203-4, Peacocke [1992]: 21, 171-5, Boghossian [2001]). If there is and cannot be any BOCHE concept, my counterexample fails.

I'll make five points in response to the claim that there are no concepts such as BOCHE. First, the move is (once again) ad hoc. Denial of BOCHE and its ilk is motivated purely by the desire to save the theory from counterexample. Second, rival views, including my preferred recognition view, can accept the existence of such defective concepts. Other things being equal, this counts in favor of the recognition view and against the inference view. (Similar points apply against the belief and intuition views). Third, the move is implausible. *Prima facie*, BOCHE is a concept. People think thoughts containing it, including THAT JERK IS BOCHE. The behavior of those who use BOCHE is repugnant precisely because they use a racist concept and express racist thoughts, not because they use no concept and fail to express any

thought.¹⁰

Fourth, one can accept the existence of “defective concepts” such as BOCHE without going “whole hog” and accepting that for any set of inference rules there is a corresponding concept. For example, one can accept the existence of BOCHE but deny that there is a concept such as Prior [1960]’s TONK. Prior describes the propositional connective TONK, whose inference rules are as follows:

$$\frac{P}{P \text{ tonk } Q} \qquad \frac{P \text{ tonk } Q}{Q}$$

TONK is obviously defective; it enables one to infer from any proposition P to any proposition Q. TONK is so defective that any agent who attempted to use it would rapidly fall into error and incoherence. It’s not implausible that the rules for TONK are crazy enough that there is no such concept. Thinking TONK-ly isn’t a possible way to think. I hypothesize that much of the resistance to accepting BOCHE as a genuine concept is based on the idea that accepting BOCHE opens the floodgates and forces one to accept all types of crazy concepts, including TONK. This is not correct. There are lots of salient differences between TONK and BOCHE. Thinkers of BOCHE, unlike those of TONK, will be led into error, but they won’t rapidly devolve into incoherence. Thinkers can navigate the world quite well using BOCHE, especially if all the Germans they interact with are cruel. Not so for users of TONK who employ the concept with any frequency. Lastly, there are considerations in favor of accepting BOCHE that do not apply to TONK. In the 1940s there appear to have been actual real world users of BOCHE. I claim

¹⁰Boghossian [2003]: 242-3 and Williamson [2009]: 140 agree.

that their behavior and psychology is better explained by typing their mental states using BOCHE than by typing by some non-defective substitute.¹¹

Fifth, and perhaps most importantly, the argument can be made without using a defective concept such as BOCHE. In the next section, I'll offer a version of the argument using the concept of the material conditional and its core inference *modus ponens*, often taken to be the paradigm on which the inference view builds its theory of all concepts.

3.5 A Counterexample Without a Defective Concept

We all encounter the concept of the material conditional \rightarrow in logic class, and think with the concept long before that. In this section, I'll argue that one can master this concept while denying that *modus ponens*, a core inference for \rightarrow , is a valid form of inference. My case is taken directly from Williamson [2008]'s case of the deviant logician (chapter 4, pp. 92-4). These are the two core inferences associated with the concept \rightarrow :

$$\frac{P \rightarrow Q \quad P}{Q} \qquad \frac{P}{R \rightarrow P}$$

Consider the deviant logician Professor X. Professor X is a distinguished professor of logic, world-renowned with many seminal publications to his name. If anyone is an expert on *modus ponens*, and on the concept \rightarrow , Professor X is. Professor X, for highly theoretical reasons, denies the validity

¹¹It's a bit controversial whether what 1940s-era agents expressed by the term 'boche', and thought with, was in fact the concept BOCHE as I've described it. It has been argued ([[Cite Ralph Wedgwood here]]) that they expressed the concept GERMAN instead. If they did not express BOCHE, this does weaken my argument, but not significantly. All that matters is that some agents could think with BOCHE, even if no one actually did so *circa* World War I.

of *modus ponens*. Perhaps he has been convinced by the alleged counterexamples of McGee [1985]. Professor X is wrong to deny its validity, but he has well-considered reasons. This thought experiment yields the following argument against the inference view.

(P1) If the inference view is true, then Professor X does not have mastery of the concept \rightarrow .

(P2) Professor X does have mastery of \rightarrow .

(C) Therefore: the inference view is false.

Premise (P1) follows from the statements of the inference view in combination with the description of the thought experiment. Premise (P2) is tremendously plausible. If anyone has mastery of \rightarrow , Professor X does. He is the one of, if not the, world's foremost experts on the conditional. The argument is valid. The conclusion follows: the inference view is false. Importantly, this argument does not rely on a defective concept such as BOCHE. The inference rules associated with \rightarrow are truth-preserving. Professor X wrongly denies the validity of *modus ponens*.

3.6 Extending the argument against the belief and intuition views

Arguments analogous to those given above can be used to defeat the belief and intuition views. The modifications required to counterexample these affiliated views are minimal. I won't go through the details.

The intuition view fares slightly better than the belief or inference views. The intuition view allows that a thinker might fail to believe a core proposition, or be disposed to infer in accordance with a core inference, yet still

master the concept nonetheless, as long as the thinker intuited, or somehow otherwise felt, that the core propositions were true. However, the case of Agnes and BOCHE still defeats the intuition view. She has no intuitions that all Germans are cruel, yet masters the concept nonetheless.

4 The Special Status Problem: Against the Sufficiency of Belief, Inference, or Intuition for concept mastery

4.1 Overview of the Special Status Problem

In this section I'll argue that the simple belief, inference, and intuition view cannot work. Believing, inferring, or intuiting are not, by themselves, sufficient for mastery. Instead, a thinker must take the core elements (propositions, inferences, etc.) to have some special status. For example, a thinker might believe core propositions because they are core or otherwise central to the concept. The initial failure of the simple belief, inference, and intuition views lends itself to an obvious fix. However, the fix must take a certain form in order to succeed. I argue that some existing ways of coping with the problem do not work.

4.2 Not Merely Belief

Consider two agents who believe all of BACHELOR's core propositions. Sofia and Felipe both believe that all bachelors are male and that all bachelors are unmarried. Sofia believes these propositions because she takes them to be central to the meaning of 'bachelor' and the concept BACHELOR. Felipe, on the other hand, believes that all bachelors are male, but that this is a matter of happenstance. Felipe is open to the possibility of female bachelors,

though he does not believe that any currently exist. Felipe's attitude toward ALL BACHELORS ARE MALE is similar to our attitude toward THERE ARE NO BACHELORS ON MARS.¹² It's true that there are no bachelors on Mars. But that has nothing in particular to do with the meaning of 'bachelor' or 'Mars' and no particular salience vis-a-vis BACHELOR or MARS. Sofia has concept mastery of BACHELOR. Felipe does not. This counterexamples the belief view. Similar counterexamples can be given to the inference and intuition views.

(P1) Felipe believes all of BACHELOR's core propositions.

(P2) If the belief view is true, then if Felipe believes all of BACHELOR's core propositions, he has mastery of BACHELOR.

(P3) Felipe does not have mastery of BACHELOR.

(C) Therefore: the belief view is false.

Premise (P1) is a stipulation of the case. Whatever the core propositions are, suppose Felipe believes them. (P2) follows from the statement of the belief view. (P3) is very plausible. Felipe is open to the possibility of female bachelors. Therefore he does not have mastery of BACHELOR. The argument is valid. The conclusion follows: the belief view is false.

¹²One might take this case as motivation to move toward a view on which the core propositions are all necessary in form. The core proposition is not "all bachelors are male" but rather "necessarily, all bachelors are male". This view does not work. There is still a difference between someone who takes "necessarily, all bachelors are male" to have special status with respect to, or to be "part of the meaning" of, BACHELOR and someone who does not. For example, x might take "necessarily, all bachelors are male" to be true because the actual world is the only possible world and actually, all bachelors are male. But x could believe that maleness has nothing to do with the meaning, or nature, of bachelors.

4.3 Sophisticated Views

One lesson of these examples is that to fully understand, and to have concept mastery, one must, in some sense, take the core truths (inferences, rules) to have some special status. This point has been recognized before (Peacocke [1992], Greenberg [2014]: 8fn.6). According to the inference view of Peacocke [1992], an agent who makes the appropriate inferences need not fully understand the concept (in Peacocke’s terminology, ‘possess the concept’ (cf. fn. 7)). To fully understand, the agent must find those inferences *primitively compelling*. Peacocke correctly recognizes the importance of the agent’s taking some special stance toward the core inferences. Peacocke’s version of the special stance is finding the inference primitively compelling.

Peacocke’s stratagem can be adopted by any of the belief, inference, or intuition views. It can be implemented in a variety of ways. That the thinker find the inferences primitively compelling is one way. One might also maintain that the thinker infer, or believe, because she takes the inferences (or propositions) to be core. Proponents of the intuition view could claim that the agent’s intuit that the propositions are core, or that taking them to be core plays some role in the creation of the intuition. There are a variety of options. I can’t and won’t canvas them all. Let’s call any version of the belief, inference, or intuition view that adopts some method of requiring that beliefs in, dispositions to infer in accordance with, or intuitions with respect to core elements a special status a *sophisticated belief/inference/intuition view*.

4.4 The All and Only Problem

Taking a special attitude toward core elements (e.g. finding primitively compelling) is still not sufficient for concept mastery. The special attitude must be taken toward *all and only* the core elements. An agent who, in effect, builds extra elements into her conception of the concept, loses her mastery of the concept, despite the fact that she bears the special attitude to every element of the core.

Consider Peacocke's theory. Suppose Malik finds the inferences from "x is a bachelor" to "x is male" and from "x is a bachelor" to "x is unmarried" primitively compelling. Suppose that these inferences completely constitute BACHELOR's core. According to the sophisticated inference view endorsed by Peacocke, Malik has mastery of BACHELOR. But, unfortunately, Malik also finds the inferences from "x is a bachelor" to "x always wears a baseball cap", "x never cleans his apartment", "x plays too many video games", and "x drinks the blood of cats" primitively compelling. Malik has an absolutely insane conception of what bachelors are like, and he takes these inferences to be part of the meaning of BACHELOR. Malik lacks mastery of the concept BACHELOR.

The obvious move is to modify Peacocke's view accordingly. On such a modified view, a thinker has mastery of C if and only if she is disposed to infer in accordance with C's core inferences and finds all and only those inferences primitively compelling. But this view is crazy. It entails that a thinker can't fully understand (master) more than one concept at a time. If an agent is to master more than one concept, she'll need to find the core inferences of at least two concepts primitively compelling. So she can't find

all and only the inferences associated with C primitively compelling.¹³ Thus she does not have mastery of C.

I see no way around this problem while sticking with the notion of “primitively compelling inference”. Analogue notions of “primitive belief” and “primitive intuition” (whatever those may be) will face the same problem. The deep problem lies in the logical form of the primitively compelling relation. It is a two-place relation: finds-primitively-compelling(thinker T, inference I). The problem can only be solved by a three-place relation. For example, one might use the “finds primitively compelling in virtue of concept C” relation: finds-primitively-compelling-in-virtue-of(thinker T, inference I, concept C). With this relation, a thinker could find some inferences primitively compelling because of the role C₁ plays in the inference, while finding other inferences primitively compelling because of the role C₂ plays. With a 3-place relation, concepts cease stepping on each other’s toes by preventing the thinker from mastering the other.

On further thought, we can motivate the need for a 3-place relation without the problem of concepts preventing mastery of other concepts. Suppose Felipe finds the inferences from “x is a bachelor” to “x is male” and “x is unmarried” primitively compelling. (Any two-place special attitude will work for the example). But Felipe finds the inference from “x is a bachelor” to “x is male” primitively compelling because of the role that MALE plays in the inference. Felipe accords this inference no special status vis-a-vis BACHELOR. Instead, Felipe has a quirky conception of MALE, according to which he finds every inference of the form “x is F” to “x is male” primitively compelling. As

¹³I make the trivial assumption that a thinker can master two concepts with non-identical cores.

a result, Felipe lacks mastery of BACHELOR, despite the fact that he finds all (and only) BACHELOR's core inference primitively compelling. The problem is that the source of Felipe's finding the inference from "x IS A BACHELOR" to "x IS MALE" is his misunderstanding of MALE rather than his understanding of BACHELOR. A 2-place relation "takes to have special status", such as "finds primitive compelling", can't distinguish between the case where a thinker accords a core element of concept C a special status because she understands C and the case where she accords special status because she misunderstands some other concept C'. Without a 3-place relation with the concept itself as a relata, a thinker could take the special attitude toward all and only the concept's core elements, but lack mastery nonetheless.

Summarizing, we've covered two distinct reasons for moving from a 2-place special attitude toward core elements to a 3-place relation. First, a 2-place relation does not allow a thinker to adopt the special attitude toward all and only the core elements of C while simultaneously having mastery of more than one concept. A 3-place relation does. Second, a thinker can adopt a special 2-place attitude to all and only the core elements of a concept C while failing to master that concept. The Felipe-BACHELOR demonstrates this point. Instead, the thinker must accord the core elements some special status with respect to C. Again, a 3-place relation yields the correct results.

4.5 Moving On

I don't think the special status problem is fatal for any of the belief, inference, or intuition view. Something more than mere belief, inference, or intuition must be appealed to, but the theories can be modified appropriately. Peacocke [1992] recognized this point. The thinker must adopt a special attitude

toward the core elements. But not even this is enough. The special attitude must be taken to core elements *vis-a-vis the concept*. The concept itself must play a role in the special attitude, whatever it is.

The belief, inference, and intuition views have problems. Neither belief, nor inference, nor intuition is either necessary or sufficient for concept mastery. The views must be modified to include a 3-place attitude the thinker can take toward a core element and a concept.

5 The Recognition View

5.1 From the Failures of Rivals to the Recognition View

In this section, I offer my own positive view of what it is required to have mastery of a concept. The view is motivated by the two failures of the inference, belief, and intuition views. These views first failed by neglecting the possibility of masters of concept C who fail to endorse core elements via belief, intuition, or a disposition to infer. Agnes has mastery of BOCHE without believing that all Germans are cruel. Neither believing in, inferring in accordance with, or intuiting the truth of, core elements is necessary for concept mastery. The second failure was ignoring the need for thinkers to adopt some special attitude toward core elements. Mastery of a concept C requires more than merely believing C's core propositions. One must believe them, and in some sense, take the propositions to have some important status *vis-a-vis* C. Neither belief in, inferring in accordance with, or intuiting the truth of, core elements is sufficient for concept mastery.

The inference, belief, and intuition views are overly demanding in one sense and not demanding enough in another. They're overly demanding be-

cause they require conceptual masters to endorse core elements. They're not demanding enough because they don't require thinkers to accord core elements a special status. The recognition view jettisons the overly demanding aspect of these views and adds the feature they neglect. According to the recognition view, concept mastery is purely a matter of recognizing core elements to have a special status with respect to the concept. That's it. Recognizing the core elements as having a special status vis-a-vis C, or to govern the meaning of C, is compatible with failing to endorse those elements.

- **The Recognition View of Concept Mastery:** An agent has mastery of concept C if and only if the agent takes all and only the elements in C's core to govern the use of C.

The recognition view is agnostic about the nature of the elements in a concept's core. The elements could be propositions, patterns of inference, mental transitions, rules, or something else altogether. The definition uses 'takes' rather than 'recognizes' because recognition is factive. If x recognizes that P, then P is the case. For this reason, it is impossible to recognize non-core elements as core. If a thinker x takes a non-core element to govern the use of C, that should count against x's mastery of C. We want the recognition view to reflect this phenomenon. Thus I use 'take' instead of 'recognize'. Recognizing a core element as governing the use of C can be factored into two components. First, the agent must take the core element to govern the use. Second, the rule must actually govern the use. The next two sections each explain one of the components.

In semantics, *meaning postulates* are axioms laid down to establish the meaning of a term. I think of core elements as playing a similar role. The

core elements help to give the concept its meaning. A different core would yield a different concept.¹⁴ Taking a core element to govern the use of a concept is similar to taking that element to be a meaning postulate that contributes to giving the concept its meaning.

5.2 Taking to Govern Use

To explain the recognition view, I must explicate what it is for an agent to *take* a core element as governing the use of C. For ease of exposition, I'll speak of core elements as *rules*. I intend to construe 'rule' as broadly as possible. On this terminology, inference patterns, propositions, and mental transitions all count as rules.

I will not offer a reductive theory of the "taking" relation. One should not view the taking relation in too heavy-weight a manner. A thinker need not have the concepts CONCEPT, GOVERN THE USE, MEANING POSTULATE, or MEANING CONSTITUTING RULE in order to recognize a core element as governing the use of the concept. Whether the thinker takes a rule to be a core rule governing use of the concept is determined by idealization from the psychology of the agent. The agent's past usage and present dispositions should be taken into account. The "takes to govern the use" relation is psychologically real. Whether an agent takes a core proposition to govern is a genuine fact about the agent's psychology. It is not merely a "stance" that

¹⁴One might wonder whether two concepts could have the same core. Probably not, for trivial reasons. C_1 's core will likely mention C_1 , not C_2 . C_2 's core will mention C_2 , not C_1 . So if C_1 and C_2 are distinct they have distinct cores. The more interesting question is whether there could be two distinct concepts with homomorphic cores, in which each element of the core of C_1 could be mapped to an element of C_2 's core by swapping occurrences of C_1 for occurrences of C_2 , and vice-versa. On this question I am officially agnostic.

theorists take toward the agent. The taking will, in general, have ramifications on the thinker's behavior. However, the best way to determine whether a thinker takes a proposition to govern the use (rather than, say, merely taking the proposition to be true) will often be by testing her behavior in counterfactual situations, idealized in various ways.

There are many ways to take a given proposition to govern the use of a concept. Consider the proposition ALL BACHELORS ARE MALE. Most straightforwardly, a sophisticated thinker might directly believe that ALL BACHELORS ARE MALE governs the use of BACHELOR. Professor X's taking of *modus ponens* to govern the use of the material conditional seems to have something like this form. More commonly, a thinker's taking ALL BACHELORS ARE MALE to govern the use of BACHELOR will be less sophisticated. She'll be confounded or perplexed when someone claims that some bachelors are not male. This perplexation will have a different feel than the perplexation she'll experience when someone claims that several bachelors live on Mars. She'll predict the interlocutor who claims that bachelors live on Mars to be misinformed. Perhaps he or she heard that the Mars rover Opportunity is a manned mission involving a crew. In contrast, she'll tend to think that the interlocutor who claims that some bachelors are not male misunderstands how to use BACHELOR and what it is to be a bachelor.

Conceivability provides a useful heuristic for gauging what an agent takes to govern the use of a concept. A proposition is *conceivable* for an agent when he or she finds the proposition possible. A proposition is *inconceivable* when she finds the proposition to be impossible. In general, an agent who takes a proposition P to govern the use of a concept will find P inconceivable.

For example, someone who takes ALL BACHELORS ARE MALE to govern the use of BACHELOR will tend to find SOME BACHELORS ARE NOT MALE impossible. Someone who does not take it to govern the use of BACHELOR will find such a situation conceivable. This heuristic yields the correct verdict in the case of Sofia and Felipe (section 4.2). Both Sofia and Felipe believe the proposition ALL BACHELORS ARE MALE. Sofia, who has mastery of BACHELOR, finds SOME BACHELORS ARE NOT MALE inconceivable. Felipe, who lacks mastery of BACHELOR, finds SOME BACHELORS ARE NOT MALE conceivable. The heuristic does not work well for deviant masters. We can flesh out the BOCHE/Agnes case (3.2) so that Agnes finds SOME BOCHES ARE NOT CRUEL conceivable, despite the fact that she takes ALL BOCHES ARE CRUEL to govern the use of BOCHE.¹⁵

In sum, I claim that there is a psychologically real difference between an agent who takes ALL BACHELORS ARE MALE to be true because it bears a special status with respect to the use of BACHELOR and the otherwise identical agent who takes ALL BACHELORS ARE MALE to merely happen to be true. In my terminology, the first takes ALL BACHELORS ARE MALE to govern the use of BACHELOR. The second does not, and thereby does not have mastery of BACHELOR. Furthermore, this distinction applies even to agents that do not have sophisticated concepts such as CONCEPT, MEANING,

¹⁵There are a variety of ways to precisify the BOCHE/Agnes case. Not all of them yield this result. In the case as originally described, it's not clear what Agnes's attitudes toward BOCHE are. There are many possibilities. Perhaps she thinks that no one is a boche. Perhaps she thinks that all Germans are boches but not all boches are cruel. Perhaps she thinks that only some Germans are boches, all of which are cruel. Or perhaps she thinks that some Germans are boches, and some boches are cruel. What Agnes finds conceivable will depend on how we precisify the case. All I claim is that on one reasonable way of precisifying the case the conceivability heuristic will not yield the correct verdict regarding Agnes's mastery of BOCHE.

or GOVERN THE USE.

5.3 Governing Use

In this section, I'll explain the sense in which a core element, or rule, governs the use of a concept. I explained earlier (section 5.1) that core elements play something like the role of meaning postulates in semantics. Meaning postulates help give a linguistic term its meaning. Similarly, core elements help give a concept its meaning. The core elements, or rules, are guidelines laid down for the use of the concept. They have normative force, indicating how one should and should not use the concept. In this sense they govern the use of the concept.

The possibility of defective concepts poses a special problem for this approach to core elements as use-governors. I've claimed that mastery of concept C entails taking C's core elements to govern the use of C. However, some concepts, such as *BOCHE*, have core inferences that do not preserve truth. (Alternatively, one could take *BOCHE* to have core propositions the conjunction of which is false). An agent who realizes that a concept's core elements do not preserve truth should in fact cease thinking (and inferring) in accordance with that element. Agnes is in that situation. She recognizes that inferring from "x is German" to "x is cruel" is a bad inference, despite the fact that *BOCHE*'s core elements suggest the inference. If core inferences need not preserve truth, in what sense do they "govern the use" of the concept?

I'd like to use an analogy with the law to explain the sense in which core elements can govern the use of a concept while failing to yield truth. Suppose that the law, and particular laws, have certain goals, purposes, or functions: to give citizens a better life, promote equality, prevent injustice, express the

will of the people, prohibit immoral action, etc. (That law, or laws, do have goals or functions in this way, that are not merely the goals of law-makers, is controversial. But the position is certainly coherent. Coherence is all my analogy requires.)

Sometimes a law achieves these goals. Sometimes it does not. When a law does not achieve its goal, it remains a law nonetheless and retains its legal normative force. The fact that the law does not, e.g., promote equality, does not thereby render the law no longer legally binding. Those who disobey the law still violate a legal norm. They can be punished. As a legal defense, “the law I broke does not achieve the goal of our legal system” will impress neither judge nor jury. If the law is immoral, then violating the legal norm may be the only way to obey a moral norm.

Rules of use have a certain function as well. They aim, broadly speaking, toward truth. They can aim at truth itself, or at the preservation of truth. (The former is more appropriate to core propositions, the latter to core inferences.) Rules may have other functions as well. But rules of use, much like the laws of a nation, can fail to achieve their aims. The core rules of defective concepts such as *BOCHE* fail in exactly this way. They fail to meet their truth-oriented functional goal. But nonetheless the rules remain rules that govern the use of *BOCHE*. They retain normative force. In some cases, the only way to obey the epistemic norm of believing truth will be to flout the meaning norm governing the use of a concept such as *BOCHE*. Agnes recognizes something like this fact about her situation. She elects to obey the epistemic norm at the cost of flouting the use norm.

In sum, there is a clear sense in which a rule can generate a norm that

governs the use of a concept. In the case of a defective concept, this norm of use might conflict with the epistemic norms of truth and truth-preservation.

6 Four Objections

6.1 Overview of the Objections

In this section, I consider four objections to the recognition view. The first says that concepts do not have cores. The second objection is that the recognition view appeals to a convoluted and under-explained notion of “taking a rule to govern the use of the concept”. The third objection is more direct. It claims that an agent can have concept mastery without taking core propositions to govern the use. Thus taking to govern use is not necessary for concept mastery. The fourth objection maintains that taking to govern use is not sufficient for concept mastery. Mastery of some concepts requires certain abilities.

6.2 Objection 1: Against Cores

Do concepts have cores? If not, then it cannot be the case that mastery of a concept is a matter of “grasping” its core. Many have thought that concepts do not have cores (Quine [1951], Williamson [2008], p.c). This is not an objection to the recognition view in particular, but rather an objection to the entire conceptual role approach to concept mastery. Behind this thought is that there is nothing “shared” by *all* those who are said to have mastered a concept.

I must admit that I have some sympathies with this line of criticism. However, I think that much of the force of the objection can be deflected by plausible softenings of the general approach. In section 2.5, I noted that,

proponents of, e.g., the belief view, could disagree about *how many* of the core propositions must be believed to qualify for mastery. For simplicity's sake, I opted to consider only views that claimed that mastery involved grasping *all* the concept's core elements. But once we relax that assumption, we can accommodate the objector's intuition that there is nothing shared by all masters of a concept. All those who fully understand JUSTICE might grasp some significant portion of the core propositions. This does not entail that there is some proposition P that they all grasp.

However, it's not implausible that, for some concepts and some core elements, all masters of the concept do share grasp of core elements. Any agent who does not recognize that arthritis involves inflammation, that swords have blades, that chairs are for sitting, or that justice requires fairness, does not fully understand ARTHRITIS, SWORD, CHAIR, or JUSTICE.

It's also important to remember that the approach to concept mastery adopted here, according to which mastery is a matter of grasping core elements (whatever "grasping" is), is not committed to the claim that a concept's core contains necessary and sufficient conditions for falling under the concept. I am tremendously skeptical of that claim. Some of the skepticism towards cores more generally (cf. Quine [1951]) arises from skepticism that concepts contain necessary and sufficient conditions for their application.¹⁶ This form of skepticism does not translate over to any of the views of concept mastery discussed here.

A core-oriented approach to concept mastery can, and I believe should, strengthen its position by broadening the notion of "core element", as dis-

¹⁶For a dissenting view, according to which users do have a grasp on necessary and sufficient conditions for the application of their concepts, cf. Peacocke [2003].

cussed in section 6.5. One can include as core elements transitions from observations or experiences to categorizations of those observations/experiences. For example one might take the transition from (i) a red experience to classification of that experience as an instance of RED (ii) an observation of Tibbles to the thought THAT IS A CAT, or (iii) a visual image of a moving horse to an ascription of GALLOPS to that horse.

Mastery of the concept will then require an appropriate relation to these core elements. A recognitional ability might constitute one way to grasp these types of core elements. Once we relax the assumption that mastery requires grasp of *all* the core elements, and broaden our conception of what cores can contain, core-oriented approaches become simultaneously more flexible and plausible.

Lastly, my task here is not to defend the idea that concepts have cores. My approach assumes that they do and that mastery involves taking some attitude towards that core. I seek the best view of concept mastery given those (admittedly controversial) assumptions. Thus, I'm somewhat happy to conditionalize my conclusion: If there is such a thing as concept mastery, and it is a matter of taking some attitude to core elements, then that attitude is recognition, not endorsement through belief, intuition, or a disposition to infer. However, I'm optimistic that the conditionalization is unnecessary.

6.3 Objection 2: Against the notion of “taking a rule to govern the use”

Opponents might object to the notion of “taking a rule to govern the use” on which the recognition view relies. Opponents might maintain that the notion is under-explained. Or they might maintain that the notion is unex-

planatory because it is too similar to the target notions of mastery and “full understanding” of a concept.

Perhaps the notion of “taking a rule to govern the use” is under-explained. I’ve done some work to get the reader on to the idea. One reason for under-explanation is that I’ve tried to be relatively agnostic about the nature of core rules. They could be propositions, rules of inference, or something else altogether. My agnosticism has a theoretical basis. I’m inclined to think that different types of concepts will require different types of core elements. Clearly, there is further work to be done in exploring the notion of “taking a rule to govern the use of a concept”. But I take myself to have made some headway.

I admit that “taking a rule to govern the use” is more similar to “fully understanding, or mastering, the concept” than believing or inferring. But I think that this is a step in the right direction. The failures of the belief and inference views are partly due to their overly reductive ambitions. To achieve a successful theory of concept mastery, we must appeal to notions that are closer in nature to understanding itself. “Recognizing/taking to govern the use” is closer to understanding. That’s part of the reason the recognition view succeeds where its rivals fail.

I also reply to the objection with the “it’s everyone’s problem” response. In section 4.2 I argued that the simple versions of the belief, inference, and intuition view do not work. There’s a significant difference in understanding between someone who merely infers and someone who infers because that inference governs the use of the concept, or plays some other special role vis-a-vis the concept. The sophisticated versions of these views accommodate the

point. But any view that accommodates the point must appeal to something like “taking a rule to govern the use”. Peacocke [1992] does so with the notion of a “primitively compelling” inference. Thus, every theory of concept mastery (or at least all the options I’ve considered) will appeal to something like “taking a rule to govern the use”. Perhaps the notion is under-explained. Or perhaps it is close to the target notion of full understanding of a concept. But everyone must appeal to some notion in the vicinity. So it’s everyone’s problem. (I admit that some “notions in the vicinity” might be more or less problematic than others. Perhaps “primitively compelling inference” is less problematic than “taking to govern the use”. I’m not sure. If it is, I’d like to see the arguments.)

6.4 Objection 3: Mastery Without Taking to Govern Use

This objection maintains that a thinker can have mastery of a concept without taking a core rule to govern the use. Thus the recognition view is false, because taking to govern the use is not necessary for concept mastery. Here’s the counterexample. Consider Professor Y, a philosopher of mind and language. She works on concepts. She has read this paper. She detests it. Professor Y believes that there is no distinction between core and non-core rules, no such thing as a rule that “governs the use” of the concept. Professor Y’s theoretical commitments lead her to deny every claim of the form “Rule R governs the use of concept C”. This thought experiment leads to the following argument against the recognition view.

(P1) Professor Y has concept mastery of many concepts.

(P2) Professor Y does not take any rules to govern the use of those

concepts.

(P3) If (P1) and (P2), the recognition view is false.

(C) Therefore: the recognition view is false.

Premise (P1) is very plausible. Everyone has mastery of at least a few concepts (e.g. AND, HERS, HERE). Furthermore, Professor Y is a world-authority on many concepts, including the concepts CONCEPT, MEANING, TRUTH, and RULE OF USE. If anyone fully understands these concepts, she does. Premise (P2) is quite plausible as well. Professor Y denies that rules govern the use of concepts. Thus she does not take any rules to govern the use of any concept. Premise (P3) follows from the statement of the recognition view, which maintains that taking rules to govern use is a necessary condition for concept mastery.

I deny premise (P2). Professor Y denies all claims of the form “rule R governs the use of concept C”. She believes many statements of the form “rule R does not govern the use of concept C”. I claim that, nonetheless, Professor Y takes many rules to govern the use of many concepts. Consider her behavior. She frowns in confusion when you talk about female bachelors. She won’t understand if you attempt to use the term ‘here’ to refer to a location other than the current one (or the location demonstrated). She’ll claim (perhaps unreflectively) that you’re incorrectly using the concept HERE. Professor Y’s belief that rules do not govern use is compatible with her taking certain rules to govern use. Furthermore, I don’t see how Professor Y could think, or use any concepts at all, if she never took any rule to govern the use of any concept. The rules guide her use of the concepts.

6.5 Objection 4: Taking to Govern Use Without Mastery

Elsewhere I have argued (citation suppressed) that mastery of phenomenal concepts such as RED_{ph} requires certain abilities, including the ability to recognize a red sensation as an instance of RED_{ph} .¹⁷ If this is correct, then it seems that someone could take whatever the core rules associated with RED_{ph} (perhaps “ red_{ph} is a sensation”) are to govern the use of RED_{ph} whilst being unable to recognize a red_{ph} experience as red_{ph} . Let Jane be a color-blind person who does not experience red_{ph} at all. She experiences $green_{ph}$ instead. She cannot recognize a red_{ph} sensation as an instance of RED_{ph} . She’s never experienced red_{ph} . But Jane, a philosopher of mind and meaning, is an expert on the semantics of ‘red’ and on the concept RED_{ph} . She certainly seems to know what the rules governing the use of the concept are. This thought experiment leads to the following argument against the recognition view.

(P1) Jane takes all and only the rules governing the use of RED_{ph} to govern the use of RED_{ph} .

(P2) If (P1), then if the recognition view is true, Jane has concept mastery of RED_{ph} .

(P3) Jane is unable to recognize a red_{ph} sensation as an instance of RED_{ph} .

¹⁷ RED_{ph} is the phenomenal, experiential, concept of red. RED_{ph} applies to experiences with a certain qualitative character (you know the one). Neither light-waves nor objects can be red_{ph} . Tomatoes, fire trucks, and strawberries are red, but they cause red_{ph} sensations.

(P4) If (P3), then Jane does not have concept mastery of RED_{ph}.

(C) Therefore: the recognition view is false.

Premise (P1) seems to follow from Jane's philosophical expertise regarding the concept RED_{ph}. (P2) follows from the statement of the recognition view. (P3) is a stipulation about the case. (P4) is a plausible claim about any phenomenal concept, including RED_{ph}. I argue for (P4) in (citation suppressed).

One simple response would be to accept the argument and weaken the recognition view by claiming that taking all and only a concept's core rules to govern use is only a necessary, but not a sufficient, condition for mastery of the concept. But doing so significantly reduces the interest and the scope of the recognition view.

I prefer to deny premise (P1). The key idea here is to consider the possibility that the ability to recognize red_{ph} experiences as instances of RED_{ph} is itself a core rule. An inference rule is a form of mental transition. The shift from an experience of red_{ph} to a labeling of that experience as RED_{ph} is also a mental transition. This transition could be a core rule governing the use of RED_{ph}. On reflection, this seems exactly correct. There is something about labeling red_{ph} experiences as red_{ph} that is central to the meaning and use of RED_{ph}. Here, we see a broadening of the kinds of elements that can occur in a concept's core.¹⁸ Before we considered propositions and inference rules. Now we include mental transitions considered more broadly. (In some sense, perhaps the transition from a red_{ph} experience to a labeling as red_{ph} is a variety of "inference", broadly construed. Little turns on the nomenclature.)

¹⁸Peacocke [1992]: 7-8 adopts a similar move.

If it is correct that the mental transition from a red_{ph} experience to a recognition that the experience falls under RED_{ph} is a core rule governing the use of RED_{ph} , then premise (P1) is false. The fact that Jane cannot recognize red_{ph} experiences as red_{ph} entails that she fails to take a core rule governing the use of RED_{ph} as such.

7 Conclusion

For now, this ends my attempt to answer the question “Under what conditions does an agent master, or fully understand, a concept?” Even if the reader does not agree with my positive theory, the recognition view, he or she should recognize that any theory of concept mastery must cope with the phenomena of incomplete understanding, deviant masters, and the special status problem. I have done my best to navigate these treacherous waters.

At root, the belief, inference, and intuition view all make the same error. They are all *endorsement views*, according to which mastery is a matter of *endorsing* elements of the core. In contrast, the recognition view is, for lack of a better term, a *recognition* view. Mastery is a form of recognizing the concept’s core. Importantly, as the phenomenon of deviant masters teaches us, recognition is compatible with a failure to endorse.

I conclude that mastering, i.e. fully understanding, a concept is not a matter of merely having certain beliefs, making certain inferences, or intuiting certain truths. Mastering a concept is a matter of recognizing, perhaps implicitly, the rules governing the use of the concept. This should not come as a surprise. Using a concept includes being subject to a normative standard associated with that concept. This normative standard is determined,

in large part, by the rules that govern the use of the concept. To understand the concept is to grasp those rules.

References

- Ball, Derek. 2013. Consciousness and Conceptual Mastery. *Mind*, **122**, 497–508.
- Bealer, George. 2002a. Intuition and the Autonomy of Philosophy. *In: DePaul & Ramsey [1998]*.
- Bealer, George. 2002b. Modal Epistemology and the Rationalist Renaissance. *Pages 71–125 of: Hawthorne, John, & Gendler, Tamar Szabo (eds), Conceivability and Possibility*. Oxford University Press.
- Boghossian, Paul. 2001. How are Objective Epistemic Reasons Possible? *Philosophical Studies*, **106**, 1–40.
- Boghossian, Paul. 2003. Blind Reasoning. *Aristotelian Society Supplementary*, **77**, 225–248.
- Burge, Tyler. 1979. Individualism and the Mental. *Midwest Studies in Philosophy*, **IV**, 73–122.
- Burge, Tyler. 1982. Other Bodies. *Pages 97–120 of: Woodfield, A. (ed), Thought and Object: Essays on Intentionality*. Clarendon Press.
- Burge, Tyler. 1986. Intellectual Norms and Foundations of Mind. *Journal of Philosophy*, **83**, 697–720.
- Burge, Tyler. 1990. Frege On Sense and Linguistic Meaning. *Pages 30–60 of: Bell, D., & Cooper, N. (eds), The Analytic Tradition*.
- DePaul, Michael, & Ramsey, William (eds). 1998. *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. Rowman and Littlefield.
- Frege, Gottlob. 1952/1892. The Thought: A Logical Enquiry. *In: Geach, Peter, & Black, Max (eds), Translations of the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell.

- Greenberg, Mark. 2001. *Thoughts Without Masters: Incomplete Understanding and the Content of Mind*. Ph.D. thesis, Oxford University.
- Greenberg, Mark. 2014. Meta-Semantics and the Distinction between Concept Possession and Concept Mastery. *Pages ??-?? of: Burgess, A., & Sherman, B. (eds), New Essays in Metasemantics*.
- Greenberg, Mark. MS. *Incomplete Understanding, Deference, and the Content of Thought*.
- Horwich, Paul. 2010. *Truth - Meaning - Reality*. Oxford University Press.
- McGee, Vann. 1985. A Counterexample to Modus Ponens. *Journal of Philosophy*, **82**, 462–471.
- Peacocke, Christopher. 1992. *A Study of Concepts*. MIT Press.
- Peacocke, Christopher. 1999. *Being Known*. Oxford University Press.
- Peacocke, Christopher. 2003. Implicit Concepts, Understanding, and Rationality.
- Peacocke, Christopher. 2008. *Truly Understood*. Oxford University Press.
- Prior, A. N. 1960. The Runabout Inference Ticket. *Analysis*, **21**, 38–39.
- Quine, W. V. O. 1951. Two Dogmas of Empiricism. *Philosophical Review*, **60**, 20–43.
- Sawyer, Sara. 2003. Conceptual Errors and Social Externalism. *The Philosophical Quarterly*, **53**, 265–273.
- Sosa, Ernest. 1996. Rational Intuition: Bealer on Its Nature and Epistemic Status. *Philosophical Studies*, **8**, 151–162.
- Sosa, Ernest. 1998. Rational Intuition: Bealer on Its Nature and Epistemic Status. *In: DePaul & Ramsey [1998]*.
- Wikforss, Asa. 2001. Social Externalism and Conceptual Errors. *The Philosophical Quarterly*, **51**, 217–231.
- Wikforss, Asa. 2004. Externalism and Incomplete Understanding. *The Philosophical Quarterly*, **54**, 287–294.

- Williamson, Timothy. 2004. Philosophical Intuitions and Scepticism about Judgement. *Dialectica*, **58**, 109–53.
- Williamson, Timothy. 2008. *The Philosophy of Philosophy*. Oxford: Blackwell.
- Williamson, Timothy. 2009. Reference, Inference, and the Semantics of Pejoratives. *Pages 137–158 of: The Philosophy of David Kaplan*.
- Williamson, Timothy. draft. How Deep is the Distinction between A Priori and A Posteriori Knowledge?