This is a preprint of an article whose final and definitive form will be published in *Erkenntnis*,

which is available online at: http://www.springerlink.com/content/102879/

H_2O, 'Water', and Transparent Reduction*

Thomas W. Polger
University of Cincinnati
thomas.polger@uc.edu
http://homepages.uc.edu/~polgertw

1. Physicalism, Reduction, and Water

Do facts about water have a priori, transparent, reductive explanations in terms of microphysics? Ned Block and Robert Stalnaker argue that they do not (B&S, 1999). David Chalmers and Frank Jackson argue that they do (C&J, 2001).[1]

The answer is important not least because some versions of the recently fashionable arguments against physicalism, of which Chalmers is the leading proponent, rely on contrasting cases like that of the identification of water and H_2O with the proposed identification of phenomenal properties and properties of brains. According to Chalmers and Jackson, it is the availability of transparent reductive explanations that secures the physicalist credentials of

* I have greatly benefited from discussions of this material with Paul Bloomfield, Tom Bontly, David Chalmers, Janice Dowell, Chris Gauker, Sandy Goldberg, Michael Lynch, Andrew Melnyk, Robert Richardson, and John Sarnecki. Special thanks go to David Henderson who commented on an early version of this paper at the Central States Philosophical Association meeting in 2005, and to an anonymously referee for this journal who provided extensive and constructive feedback. I am grateful for feedback from the audiences at the CSPA and at the University of Toledo. Work on this project was supported in part by a fellowship from the Charles P. Taft Research Center at the University of Cincinnati.

[1] See also Jackson (1982, 1998) and Chalmers (1996).

water.[2] For Chalmers (1996) it is the impossibility of such a reduction for conscious mental states that demonstrates that physicalism is false as a general thesis. Jackson (1998) rejects the contrast; but according to him it is the requirement for transparent reduction of both water and consciousness that creates the physicalists' need for conceptual analysis. So humble water turns out to be rather important. If no transparent reductive explanation is available for something like water, whose physicalist credentials are uncontested, then Chalmers is denied one of his key arguments against physicalism and Jackson is denied his argument that physicalism requires conceptual analysis.

In this paper I argue that C&J's positive arguments for the in-principle availability of a priori, transparent, reductive explanations in the cases of ordinary macroscopic facts are ineffective. My objection develops some remarks made by B&S, and supports their critique. Moreover, I show that C&J's purported repudiation of the B&S objection is inadequate. I conclude that the case of water and H_2O stands as a counterexample to the view that physicalism requires a priori, transparent, reductive explanations of all facts in terms of strictly physical facts. Because such reduction is not required, C&J have not secured a crucial premise in their argument that analytic physicalism (what Chalmers calls "Type-A physicalism") is the only legitimate sort. This conclusion has consequences for C&J's defense of conceptual analysis as an essential philosophical method.

I am certainly not the first person to argue that one or more of Chalmers's and Jackson's arguments fail.[3] Yet as C&J have advanced (individually and jointly) a growing body of work, there continues to be substantive disagreement about just what their arguments are. Because the

[2] For convenience, I will often abbreviate "a priori, transparent, conceptual reduction" by "transparent reduction." Sometimes I will speak of "a priori reduction" or "conceptual reduction" to emphasize those aspects of the C&J notion. I will use the full expression when there is a risk of confusion.

[3] Some that have been particularly helpful to my own thinking are Byrne1999, Hawthorne 2002, Hill and McLaughlin 1999, Lycan 2003 and forthcoming, McLaughlin 2005, Wilson 2005, Wright 2007, and Yablo 1999.

arguments are sometimes elaborate and are often cast in terms of some specialized philosophical devices, one sometimes gets the sense that both the problems and solutions will be "technical" and therefore best left to those with a taste for such things. Against this background, one of the contributions of the present paper is to identify an overlooked line of inquiry that does not depend on either accepting or rejecting C&J's technical apparatus. This is possible because the technical apparatus (e.g., the appeal to two-dimensional modal semantics) plays an only secondary role in C&J's core argument—a point on which they themselves insist (2001).[4]

I should also say up front that for the purposes of this paper I concentrate on the version of Chalmers' and Jackson's reasoning expounded in their joint 2001 paper. This paper is presented as a self contained exposition of both their positive view (2001: §§1-4) and their response to Block's and Stalnaker's objections (2001: §§5-8). Although they provide a number of other arguments in various works (especially Chalmers 1996 and Jackson 1998), I presume that the 2001 paper includes the resources that C&J take to be sufficient for deflecting the Block and Stalnaker critique.[5] There may be other and independent arguments against physicalism "in the neighborhood," but those are not my present concern.

Finally, it goes without saying that a refutation of C&J's arguments against physicalism does not establish that physicalism is true. But it may help us to get a better handle on the physicalist thesis.

[4] In some sense I view the apparatus as a red herring, distracting us from the main argument. As will be clear, the two-dimensional semantics at most plays the role of defending some premises and formulating some replies to critics. If my assessments hold up, then those replies to critics turn out to be question-begging. (See §§3-4.)
[5] I sometimes appeal to other works to support my interpretation of C&J 2001. But those appeals should play no direct role in my direct line of argument.

2. How to Reductively Explain Water in Microphysical Terms

In the next section we will assess the disagreement between C&J and B&S. But first we need to understand C&J's model of reductive explanation and the reasoning behind it. Because Chalmers' and Jackson's views are often misunderstood and because my critique will depend on being very clear about just exactly what they claim and what arguments they take to support their claims, it will be worthwhile to lay out their line of reasoning meticulously.

Begin with what is not disputed: Nobody has ever actually provided a complete reduction of water to H_2O, much less to physics proper. Neither C&J nor B&S suppose that anyone has.[6]

But according to C&J an explanation of non-basic water facts in terms of strictly physical base facts is in-principle available. And they provide a recipe for how it could be achieved. The central idea behind C&J's account of transparent reductive explanation is that an ideal reasoner who knows all the relevant physics facts and who has the concept 'water' would be in a position to judge that if the base facts are a certain way then the extension of 'water' is H_2O, and thereby be in a position to know various facts about water.[7] The ideal reasoner would not need any additional information to draw this conclusion, and in that way the conclusion may be thought of as a priori "derivable" from the base facts, given the concepts.

What do C&J have in mind when they claim that the ideal reasoner has the concept 'water'? Having a concept is to be thought of as knowing (perhaps implicitly) the truth of certain conditionals. As a very rough gloss, we can say that the ideal reasoner knows a priori something like the conditional W:

---

[6] In fact, it is part of C&J's argument that no explicit and finitely statable reductive explanation may be available.

[7] Such facts would include, among others, the fact that water boils at 100°C, that the oceans are filled with (impure) water, and that the substance kind water is identical to the substance kind H_2O. Although I am wary of "fact" talk, I'm not going to fuss over it here. I am simply adopting the "fact" talk from C&J and many of their critics, and it is a convenient device for presenting my argument. I take it that we can restate the matter in terms of properties, sentences, or propositions if we care to do so—though I don't suppose that the translation would be trivial.

(W) If H_2O is the waterish stuff around here (i.e., plays the water role) then water=H_2O.

It may be that having a concept (e.g., 'water') is to be identified with the knowledge of certain conditionals (such as W and its kin); or it may be that having a concept simply gives one such knowledge. The answer is not crucial to the current discussion. What is important is that one who has a concept *ipso facto* knows certain conditionals. Knowledge of these conditionals is, in that sense, conceptual knowledge.

As written, W may appear to be trivial; but this appearance is misleading. By "waterish" and "the water role" we abbreviate some strictly physical specification of the characteristics of water. I will follow C&J in speaking of the "waterish" properties of H_2O, but it is important that this way of talking not unduly influence our intuitions about whether W is knowable a priori.[8] Likewise, I will follow C&J in pretending that facts about H_2O are genuinely microphysical facts. Thus I'll for now set aside questions about the physical credentials of chemical entities and properties, as well as questions about how to specify those properties and entities that are strictly physical. I shall assume, with C&J, that those questions can be adequately addressed. With these qualifications in place, conditional W may be supposed to contain only one reference to non-physics facts, viz., the reference to water in the consequent.

Now consider what can be known by an ideal reasoner who knows all the strictly physical facts about the world and who has the concept 'water'. For the ideal reasoner transparent reductive physical explanation of water proceeds as follows:

(1) H_2O is the waterish stuff around here. (empirical, with indexical grounding)

[8] I'll continue to talk about "waterish properties and the "water role." But it might be best to replace (W) with (W*) which uses the abbreviation Ω for the microphysical specification:

(W*) If H_2O is the Ω stuff around here then water=H_2O.

The reader is invited to make this substitution throughout if it will be helpful.

(2) If H_2O is the waterish stuff around here then water=H_2O. (W)

(3) Water=H_2O. (1, 2, MP)

According to C&J, the strictly physical base facts *imply* the non-basic water facts in the sense that it is not possible for the base facts to be true and the water facts false. The concept 'water' combined with the physical base facts (1), put the ideal reasoner in a position to recognize ("deduce") facts about water (3).[9] This is because the concept 'water' gives one tacit knowledge of conditionals like W—as evidenced by the ideal reasoner's ability to "deduce" (3) from (1). As C&J put it:

> if a subject possesses the concept 'water', then sufficient information about the distribution, behavior, and appearance of clusters of H_2O molecules enables the subject to know that water is H_2O, to know where water is and is not, and so on. This conditional knowledge requires only possession of the concept and rational reflection, and so requires no further a posteriori knowledge. (323)

On the view advanced by C&J, it is this a priori justified "conditional ability to identify [a] concept's extension" (C&J, 324, italics removed) that underwrites the possibility of transparent reductive explanation. Facts about water are "reduced" to facts about H_2O precisely because the ideal reasoner can see that the H_2O facts imply the water facts. This knowledge is glossed by premise (2).

---

[9] A parallel "explanation" can be given using a different conditional that is said to be known by one who has the concept 'water':

(1) H_2O is the waterish stuff around here. (empirical, with indexical grounding)

(2) If H_2O is the waterish stuff around here then water boils at 100°C. (W**)

(3) Water boils at 100°C. (1, 2, MP)

This version may be preferred by those who have the intuition that water = H_2O is already a conceptual truth in its own right, in virtue of being an identity claim.

The above line of reasoning counts as *explaining* the water facts in terms of the base facts, because the ideal reasoner recognizes that the base facts necessitate the water facts by implying them. This connection is self-evident (transparent) to anyone who has the conditional ability—who has mastered the concept 'water' and thereby knows (something glossed by) the a priori conditional W. So the conditional ability to determine the extension of 'water' in principle allows one to transparently and reductively explain water facts in terms of H_2O facts using only the concept 'water' and the distribution of basic physical stuff in the world (and perhaps some other neutral or semantic facts.)[10]

Moreover, according to C&J anything short of this kind of transparent reductive explanation—e.g., "reduction" by a posteriori identification—would carry less than conceptual or logical necessity (implication), and would therefore not be an "epistemic reduction." Any weaker connection than logical determination leaves open the possibility that water and H_2O are only contingently or nomologically correlated, accidentally coinciding in the actual world. In the language that Chalmers adopts from Levine (1983), it would leave an "explanatory gap" between water and H_2O. Only "epistemic reduction" can show how the H_2O facts determine the water facts. So any successful epistemically transparent reductive explanation must have a tacit a priori conceptual component, which can be known by conceptual analysis of the reduced term.

It is important that the ideal reasoner knows something that can be glossed by W, but need not know any explicit definition or analysis of 'water' in order to count as having knowledge of the extension of the concept. The kind of "conceptual analysis" endorsed by C&J is therefore not as strict as the traditional variety that aimed at necessary and sufficient

[10] I am not presently concerned about the need for physical closure, indexical facts, or facts about experiences, which is acknowledged by C&J (2001). I will nevertheless periodically remind us that these additional facts are needed by C&J, so we don't forget that they do not come for free. (See Lycan 2003.)

conditions. C&J deny that conceptual analysis requires discovering necessary and sufficient conditions. They thereby concede that there may be no finitely statable reductive explanations. C&J they support their view of a priori conceptual analysis by analogy with the famous Gettier (1963) discussion of knowledge. If G is the complete description of a some epistemic scenario like a Gettier story, then an ideal reasoner who has the concept 'knowledge' will be a position to know whether that concept applies to the justified belief in the scenario. If G is a Gettier story then the ideal reasoner will recognize (and in that sense, "deduce") that the justified belief does not count as knowledge. That is, from G (the Gettier scenario) the ideal reasoner "deduces" the conclusion that Smith does not have knowledge. According to C&J, this shows that the ideal reasoner implicitly knows something that can be roughly glossed by K:

(K) If G obtains, then Smith's justified true belief is not knowledge.

Thus the valid deduction is completed:

(4) G, the Gettier scenario obtains. (empirical, hypothetical)

(5) If G obtains, then Smith's justified true belief is not knowledge. (K)

(6) Smith does not have knowledge; justified true belief is not sufficient for knowledge.

But what justifies the alleged knowledge of K? According to C&J, knowledge of K comes from mastery of the concept 'knowledge' and nothing more. Since it is plausible that knowing K does not require knowing any explicit definition or analysis of 'knowledge', C&J conclude that a priori analysis does not require knowing any such explicit analysis. Nevertheless, K itself is a bit of conceptual knowledge has in virtue of mastering the concept 'knowledge', as revealed by one's ability to make armchair judgments about epistemic scenarios.

Assuming that the Gettier reasoning is an example of conceptual analysis, C&J suppose that the analysis of water is analogous. In the Gettier case, the base facts are facts about Smith's

beliefs, their epistemic merits, and the arrangement of stuff in the world. We are there concerned with whether the non-knowledge base facts are sufficient for determining the distribution of knowledge in the world. It appears that they are: The facts as presented determine that Smith does not have knowledge in situation G. As with the water case, the Gettier reasoning is said to count as conceptual analysis (rather than, say, ordinary theorizing) because it does not depend on any obviously a posteriori facts. The only empirical information required appears in the antecedent of the conditional, so it plays only a hypothetical role in the analysis. What can be known a priori is that if the antecedent is true then the concept applies or does not; and that does not require knowing the truth value of the empirical antecedent. If one understands the concept 'knowledge' then one can know K a priori. Yet from K and the truth value of its antecedent, one can recognize ("deduce") with no additional empirical information ("a priori") that Smith's justified true belief is not knowledge: "Gettier's argument was an a priori argument, in which empirical information played no essential role, and its conclusion is a paradigmatic example of a non-obvious a priori truth" (C&J, 321).

There are two noteworthy features of the sketch of reductive explanation by conceptual analysis that is offered by C&J. First, if C&J succeed in showing that the case of water and H_2O fits their model, they still have not have shown that this is the only model of reductive explanation or that it is the model that physicalists must adopt. At best they will have provided sufficient conditions for one kind of reductive explanation.

Second, even if they convince us that water has been "epistemically reduced" to H_2O, they will not have demonstrated that the "reduction" fits the model that they outline. There may be other kinds of epistemically reductive explanation that account for the reduction of water, other ways of closing the explanatory gap. C&J cannot rule that out because they provide only a

model of reductive explanation but no complete reduction of water to H_2O. Here I am merely holding C&J to the same standard to which they hold their critics. They rightly argue that just because macroscopic or phenomenal facts may be "reducible" in some sense, it does not follow that they are "epistemically reducible" in the way demanded by their account (2001: 356). By the same reasoning, demonstrating that some fact is "epistemically reducible" (that there is no explanatory gap for that fact) does not establish that every successful "epistemic reduction" follows the C&J model of transparent a priori conceptual analysis. So even if "epistemic reduction" is required of physicalists, C&J have to demonstrate that their model is the only reductive model that would yield "epistemic reduction." To secure their conclusions, C&J need an independent argument to rule out alternative accounts of the explanatory success. It is important that the argument be independent because insofar as arguments against the alternatives presuppose the positive sketch, they will be question-begging.

The bottom line is that it is not enough for C&J to deflect objections to their account by arguing that the critics do not satisfy their (C&J's) reductive demands. C&J must argue that the critics cannot be correct—that there is no other satisfactory model of physically acceptable reductive explanation. This is a burden that C&J accept. Responding to B&S, C&J offer a rebuttal that takes the form of a *reductio ad absurdem*. It is this argument that we must now scrutinize.

3. Block and Stalnaker's Objection, and the Reductio Response

Block and Stalnaker (1999) develop a two part argument against the claim that reductive explanation relies on a priori conceptual knowledge. The first part of the argument concludes that no explicit analyses of macro-concepts (like 'water') are available to a priori justify the

alleged conditional ability to apply the concepts. The second part of the argument acknowledges that C&J do not rely on explicit analyses for their sort of a priori conceptual reduction. But, B&S argue, transparent reduction of macro-facts fares no better on the view favored by Chalmers and Jackson (together and individually) according to which the contents of concepts are functions from worlds to extensions, rather than being like descriptions or definitions. The B&S objection is important because both sides agree that physicalism is compatible with ordinary and scientific facts about water. If water is not reducible in the way that C&J suppose, then either C&J are wrong about what is required for reduction or else physicalism does not require reduction at all. And, of course, in order to defeat B&S's objection C&J will have to provide responses that do not depend on the correctness of their account of reductive physicalism, or face the charge of question-begging.[11]

The central element of C&J's response to B&S is a reductio argument. They briefly describe the B&S objections, very much as I did above, as attacking the C&J model of transparent reductive explanation: "The intended upshot is that the existence of these a priori entailments is doubtful, and that reductive explanation does not require a priori entailment" (C&J: 337). The a priori entailments, recall, are the sort that allow the ideal reasoner to recognize

[11] The details of the B&S argument will not matter for the present considerations. Such an assertion rightly raises suspicion. To see why it is agreeable in this case, we must step back and observe the overall structure of the dialectic. First, C&J advanced a view about physicalism and transparent reductive explanation. Then, B&S offered an objection in two parts. In response, C&J themselves employ a two stage strategy. First, they argue that B&S's critique cannot be correct (C&J, 2001, §5.1: 336-338). This stage is a reductio argument that does not even mention the details of B&S's argument. Second, supposing that they have defused the alleged counterexamples by way of the reductio, C&J argue that their favored account can handle the specific claims advanced by B&S (C&J, 2001, §5.2-5.6: 338-350). Here C&J do engage in the details of the B&S argument. But because the second stage of the defense presupposes that the initial reductio stage was successful, the whole defense hangs on that reductio. If the first stage fails, then the B&S counterexamples stand and C&J's second stage arguments are question-begging. So if we confine our attention to C&J's initial reductio argument, the details of B&S's discussion will not yet come into play. That is my strategy, and for that reason my discussion can safely omit the details of B&S's arguments. As is characteristic of the target of reductio arguments, what is important to C&J is the conclusion that B&S draw rather than the details of their route to it.

("deduce") whether something is an example of water or knowledge, given some base empirical facts and the concepts 'water' or 'knowledge'. But, C&J argue, B&S cannot be correct:

> An argument with [the structure of the B&S arguments] cannot successfully make a case against a priori entailments. To see this, it suffices to note that if this sort of argument succeeds, it succeeds equally in making a case against the a priori Gettier entailments discussed earlier. It is at least as plausible for 'knowledge' as for 'water' and for 'life' that there is no explicit analysis to support the entailments. And the general criticisms of this use of the two-dimensional framework presumably apply equally to its use in the case of 'knowledge'. But nevertheless, the a priori Gettier entailments discussed earlier exist; or at least, it is clear that this sort of argument does little to make a case against them. So by parity, this sort of argument does little to make a case against the a priori entailments we are concerned with. (C&J: 337)

Even as reductio arguments go, this one is particularly bold. The reasoning is that if B&S were correct, then their arguments would show that we do not have a priori conceptual knowledge in the case of the Gettier argument; but everyone agrees that we do have such knowledge, therefore B&S—and any other arguments of "this sort"—must be mistaken. It is clear that the "sort" of objections that this reductio is supposed to undercut are any that claim that conditionals like W are not a priori conceptual truths. So C&J's response is intended to have the generality that they need in order to show that any successful gap-closing explanations must have an a priori conceptual component.

A striking feature of C&J's reductio response is that entirely eschews their formal apparatus of explaining modal intuitions in terms of conceivability using primary and secondary

intensions. All the weight is borne by the analogy between the water case and the Gettier case, and by the assertion that the Gettier case is an example of a priori conceptual analysis that fits the C&J model.[12]

Of course C&J subsequently have much to say about the specific proposals of B&S. The second stage of their paper (§5.2-5.6: 338-350) concentrates on responses that C&J describe as "arguments that go beyond what is suggested by this [reductio] structure" (2001: 338). The second stage response should be seen as an exploration of which of B&S's assumptions lead to the alleged reductio that we do not have a priori conceptual knowledge about 'knowledge' when reasoning about the Gettier scenario. Importantly, as we would expect of follow-up arguments to a reductio, these second-stage critiques are not independent of the reductio itself. They suppose that the reductio conclusion is false; and in this case that comes to reasserting that we can have a priori conceptual knowledge of certain conditionals concerning, e.g., water or knowledge. That is, they assume the success of the reductio response against B&S's counterexamples. So the reductio response (C&J, 2001, §5.1: 336-338) is prior to and independent of C&J's second-stage discussion. If the reductio can be dismantled, then C&J's whole stable of responses to B&S is undermined. This is my strategy.

4. Concepts, A Prioricity, and the Reductio Response

C&J's reductio invites the critic to accept a dilemma: Either argue that we do not in fact have knowledge about knowledge, of the sort supposedly on display when we make Gettier judgments; or else argue that there is a disanalogy between the water and knowledge cases that

[12] Strictly speaking, the conclusion is that either there is a priori conceptual knowledge in the Gettier case, or at any rate "it is clear" that the B&S style argument has no tendency to show that there is not. That is not so clear to me; and I am puzzled as to what argument C&J have in mind. I will focus on the first possibly as the main objection, which I take it is their leverage in the reductio.

breaks the parity of reasoning.[13] But there is a third option. We may accept that the analogy is apt but argue that both cases work differently than C&J suppose. I will pursue the third option. Granting for the sake of argument that the cases are analogous and that we do have knowledge in the Gettier case, there is still reason to doubt that the knowledge we have is a priori conceptual knowledge.

In fact B&S have suggested a reason for questioning C&J's presentation of the analogy between the water and Gettier examples. But B&S do not pursue it as their main line of critique and C&J consequently fail to appreciate its force. The basic idea is that it may be that not all armchair philosophical knowledge is *conceptual* knowledge. We should ask, what reason is there to suppose that the ability to make certain judgments about knowledge or water is justified *solely* by the empirical distribution of basic stuff and our philosophical or macroscopic concepts? Why think that conditionals like W or K are logical or conceptual truths? W has the flavor of a tautology due to the occurrence of "waterish" and "the water role" in the antecedent. But that should not mislead us; we must remember that those expressions are just stand-ins for a detailed physical story that makes no mention of 'water'. Why suppose that W is an a priori or conceptual truth, given that it is not a tautology?

C&J encourage us to think of W or K as conceptual truths by speaking of the ideal reasoner's judgments about hypothetical situations as "deductions." But at best they are entitled to the claim that the ideal reasoner has the ability to make certain judgments or recognize certain truths. We have been offered no independent reason to think that those recognitional capacities must be based on conceptual mastery alone. As William Lycan writes, "All parties to the present

[13] C&J acknowledge that there are some differences between the cases because the water case involves natural kinds and the knowledge case may not, but I will not pursue that difference here. And for my part, I don't know whether or not we *know* that Smith does not have knowledge, in the Gettier scenario. So I will set these concerns aside for now.

dispute agree that the ordinary-macroscopic supervenes on the microscopic. It simply does not follow that a physically identical world in which the higher-level facts do not obtain is 'logically' or a priori impossible as opposed to merely impossible" (Lycan, 2003; see also, Lycan, forthcoming). What Lycan suggests here is that knowledge of what is "merely impossible" may not be a priori or conceptual. For all that C&J have argued, our judgments about knowledge and water can be based on substantial philosophical intuitions or tacitly held empirical theories, rather than semantic or conceptual knowledge about the extensions of concepts.[14]

The suggestion that I am developing is briefly raised in a passage from B&S. Regarding the possibility of knowing whether water is necessarily a liquid, they write:

> This seems to be armchair reasoning, reflection that does not include any obvious reference to real experiments, so it is tempting to conclude that this reflection just unfolds our concepts in a totally a priori way. But what this conclusion misses is that our reasoning about the proper epistemic response in various counterfactual situations is informed not only by our concepts, but by implicit and explicit theories and general methodological principles that we have absorbed through our scientific culture—everything that the "we" who are performing these thought experiments believe. What people should rationally say in response to various

[14] Lycan is even more permissive than B&S and for the sake of argument I am prepared to go along. We can allow that our knowledge of W or K could even be a priori and yet fail to be conceptual knowledge—thus not dependent on any conceptual analysis, explicit or otherwise. I am willing to allow C&J room to seize a solution involving what Horgan and Henderson (2002) call "high grade" a priori knowledge. This, I take it, would involve substantial synthetic a priori knowledge. I can leave this option open to C&J because it would carry a very high cost for them to accept. If that is C&J's position, then it has indeed been misunderstood and the whole debate can be focused on such a substantial philosophical commitment that is not grounded in semantic or conceptual concerns.

Although for the sake of argument I could allow that philosophical or theoretical judgments may yet count as a priori, in fact it is quite controversial whether there is any such purely a priori and conceptual knowledge. If we are at all inclined towards Quinean skepticism about the a priori then we will be rather dubious. But a prioricity is not my central focus in this paper.

> hypothesized discoveries will vary depending on their experience, commitments and epistemic priorities. (B&S, 43)

B&S immediately back off from the direct attack on the claim that armchair knowledge is a priori and conceptual. Instead they pursue the angle that *whatever* the nature of the connection between water and H_2O the physicalist may be in a position to maintain that the same relation holds in the controversial case of phenomenal properties. I agree. But the objection that I have in mind puts emphasis on a different part of B&S's observation.

There are at least two criticisms implicit in the above quoted passage from B&S. The first, which is their primary concern and is the aspect to which C&J respond, is that armchair reasoning may not be a priori reasoning after all. This is because "our reasoning about the proper epistemic response in various counterfactual situations is informed… by implicit and explicit theories and general methodological principles that we have absorbed through our scientific culture—everything that the 'we' who are performing these thought experiments believe." If some of those beliefs are a posteriori, then the whole line of reasoning is as well.

The second criticism, which B&S do not pursue and C&J do not appreciate, is that armchair reasoning (whether a priori or not) may involve reflection that "is informed not only by our concepts" and that goes beyond "just unfold[ing] our concepts." Here the emphasis is not on a prioricity, but rather on whether armchair reasoning invokes philosophical or theoretical resources other than those explicitly or implicitly encoded in our concepts. This latter point is my focus. Our armchair reasoning may, for all I know, involve implicit knowledge of certain conditionals, or the having of some conditional abilities. Similarly, Terry Horgan and David Henderson observe that, "what one learns *while* acquiring a concept is not what one learns *in* or *by virtue of* acquiring that concept. From the start, then, one acquires a mix of conceptually-

based and non-conceptually-based information and transitions" (2002). The question is: What reason do we have for supposing that the kind of reasoning that C&J describe "just unfolds our concepts"? Why suppose, that is, that when making judgments about hypothetical cases the ideal reasoner is merely examining her concepts rather than reasoning about the world?

One might think that the very fact (if it is a fact) that there is no explanatory gap in the case of water and H_2O shows that water facts are conceptually entailed by H_2O facts. But that would be to beg the question. The present suggestion is precisely that closing explanatory gaps does not require reductions involving conceptual entailment. Likewise, C&J cannot appeal to the fact that competent users of the language or concepts are able to make correct judgments about hypothetical situations, e.g., about water or knowledge, for that ability is granted by both sides in the present debate. The question at hand is whether that ability (assuming that we have it) is merely a matter of unpacking one's concepts to determine their extension—an exercise of what we might call *semantic* intuitions—or whether it engages additional resources—such as distinctly *epistemological* or *metaphysical* intuitions or theories.[15]

In essence, the issue at hand is the metaphilosophical question of whether all philosophical "analysis" involving hypothetical cases is "conceptual analysis." B&S, Lycan, and I are arguing that it is not. What is needed from C&J is an argument that there is no non-physics knowledge about ordinary objects other than what is conceptually entailed by physicsknowledge. That seems to be their position. But it is not a view that C&J can take for granted, or that they can assume any physicalist must endorse.

Now where do things stand? To answer B&S, C&J must defend their model without relying on their own preferred framework, for it is the argument for that framework that is in

[15] By "semantic"or "conceptual" intuitions, here, I refer to whatever abilities to make judgments about hypothetical cases that are conferred solely by mastering the meanings of terms or concepts.

question. The only such argument that they offer is the reductio response. But that is not convincing because C&J have not independently established that the benchmark Gettier case involves only conceptual knowledge rather than tacit philosophical or theoretical knowledge. More generally, C&J have not established that the ideal reasoner's correct judgments about actual and hypothetical cases engage only conceptual knowledge. If this is correct, then C&J have given us no reason to suppose that water is transparently, reductively explainable in the way that they demand. And if water is not, then we shouldn't expect that anything else is—least of all consciousness.

5. Conceptual Truths and Superdupervenience

I do not deny that we may have armchair philosophical or theoretical knowledge that is revealed in our judgments about hypothetical cases, as involving knowledge or water. On the proposal under consideration, which C&J have not ruled out, such knowledge may have sources other than the meanings of our concepts. Philosophical or theoretical knowledge may be epistemological, metaphysical, or scientific in nature, rather than semantic or conceptual. We may have philosophical or theoretical knowledge (e.g., about knowledge and water) that goes beyond knowing the intensions of the concepts 'knowledge' or 'water'. This view is not unprecedented, as I will elaborate in the next section. But first a few words of clarification are in order.

I acknowledge that for C&J the idea of metaphysically necessary truths that are not basically conceptual is nonsense. They claim that there is no reason to believe that there is any metaphysically necessary truth that is not conceptually necessary truth. Such truths would involve a distinctively metaphysical modality—"strong" necessary connections that are neither

logical nor merely nomological. And C&J deny that there is any reason to believe in "strong" necessity.[16] The C&J view has a distinguished history among various kinds of empiricists, logical positivists, ordinary language philosophers, Wittgensteinians, and perhaps going back to Hume. I have done nothing thus far to argue that the view is incorrect. I have simply helped myself to some resources that prima facie involve its denial. This is not question begging on my part because while there is no doubt that many physicalists have denied that there is a distinctively metaphysical modality, many others accept that there is a metaphysical modality and the burden is on C&J to show that this is incompatible with physicalism.

Second, C&J portray their view of a priori transparent reduction as a simple consequence of the physicalists' commitment to the thesis that all facts supervene on base level facts. If water facts supervene on H_2O facts (or H_2O facts plus some neutral non-water facts), then facts about H_2O "entail" (in the sense of determination or necessitation) facts about water. And if that entailment or determination can be recognized by an ideal reasoner who knows the base facts and has the concept 'water', then C&J say that the water facts are "implied" by the H_2O facts.[17] This pattern of reasoning makes it seem as though anyone who accepts the idea that macro-facts supervene on lower-level facts is thereby committed to C&J's view of implication and reduction. Because physically-based supervenience is usually thought to be at least a necessary condition on physicalism, the physicalist seems obligated to endorse this framework.

---

[16] It is important to note that neither Chalmers (1996) nor Jackson (1998) provide a positive argument for the view,; rather, as in C&J 2001, they try to undermine arguments for the alternative. Even if successful, this would not *ipso facto* demonstrate that there are reasons to adopt their preferred view. (Byrne 1999, Hill and McLaughlin 1999, Lycan forthcoming).

[17] Specifically, if B-facts supervene on A-facts, then "if A then B" is true, and C&J say that the A-facts *entail* the B-facts. When this "entailment" (i.e., conditional) is justified a priori, then C&J say that the A-facts *imply* the B-facts. (For C&J, the conditional is justified a priori if an ideal reasoner who knows the A-facts and has the B-concepts can correctly judge that the conditional is true.)

This way of presenting C&J's reasoning may seem compelling, but it disguises the radical parts of the proposal. In particular, as I have been at pains to point out, C&J assume that the ideal reasoner's recognitional ability must be grounded solely in her mastery of concepts, and likewise that mastery of concepts alone would allow the ideal reasoner to make justified judgments about various hypothetical situations. It is this logical or conceptual aspect of their account that I have been calling into question. According to the proposal I am entertaining, the ideal reasoner might employ tacit philosophical or theoretical knowledge rather than solely logical or conceptual knowledge. If there is such extra-conceptual philosophical or theoretical knowledge, then of course we shall want to know much more about it. (A topic to which I will turn in §6.) But at present the work we require can be done by the mere possibility of such knowledge, whatever its nature might be. And C&J have done nothing to rule out this possibility.

Now it may be that C&J only care that conditionals like W and K are known a priori. That is, perhaps the conditional knowledge about water or knowledge could fail to be conceptual and still be a priori. If so, then W and K would be examples of substantive synthetic a priori knowledge. This is probably more rationalism than C&J want to take on board. But it is open to them to do so. However, if they avail themselves of this option they undercut much of the motivation and rhetorical force of their conceptual approach. Is it really plausible that physicalists are ipso facto committed to such a strong form of rationalism? This will seem, to many, to be a reductio of the C&J strategy. Something must be wrong with their "serious metaphysics" if it has this consequence.

Finally, it should be clear that the question of whether facts about water or about consciousness can be a priori, transparently, reductively explained is a question about the resources to which a physicalist is entitled. In Jackson's (1998) way of posing the problem, the

issue is where to "locate" the B-facts (about water, knowledge, consciousness, etc.) among the base A-facts (the strictly physical facts.) The question of where to locate the B-facts comes down to questions about the available "space" and about what counts successfully "locating" a fact. According to C&J, the physicalist is entitled only to the A-facts of physics (distribution of matter and energy, and laws), along with supposedly neutral facts of logic, mathematics, and semantics (logical, analytic and conceptual truths, broadly speaking), and facts that are a priori entailed by the physical and broadly logical/conceptual facts—i.e., that are reductively, transparently explained by the physical facts. C&J also admit the need for indexical or index-fixing facts, and a closure clause that says there are no other basic facts.[18] These are the only acceptable "locating" resources. On this view, the only way of locating facts that are not properly facts of physics is to show that they are a priori conceptual facts or are "implied" by and transparently reducible to the facts of physics. This highlights the importance to C&J of insisting that facts that we all accept as physically unimpeachable, such as facts about water, can be located by transparent reduction. Facts about water are not part of fundamental physics, and are not merely semantic. As they are uncontested physically acceptable facts, water facts had best be a priori, conceptually, transparently reducible. If, as I claim, some ordinary macro-physical facts are not transparently reducible, then C&J must have misunderstood physicalism, reduction, or both.

The issue of the acceptability to physicalists of explanatory resources that are not themselves properly part of physics has long been pressed.[19] The way of stating the problem that I find most lucid, and after which my own discussion is modeled, is that of Lynch and Glasgow

[18] These additional facts are problematic, but I will not pursue that issue herein.

[19] See, e.g., Horgan 1984, 1993. The versions by Hellman and Thompson (1975) and Crane and Mellor (1990) are also influential.

(2003).[20] Assume that some B-facts supervene on some A-facts. (For example, the A-facts may be facts of physics and the B-facts may be non-physics facts, e.g., facts about water or knowledge. But the picture is supposed to be quite general.) If so, then there will be facts about the supervenience relations, S-facts. These will be the connecting facts that allow us to "locate" B-facts among A-facts. Now what is the status of the S-facts? They are facts about the connection between A-facts and B-facts, so they cannot themselves be either A-facts or B-facts. On pain of regress, they cannot supervene on A-facts or B-facts, for then we would have S*-facts (about the first supervenience relation) to account for, and so on. So, stipulating that "broadly logical" (Chalmers 1996) facts are A-acceptable, it looks as though the supervenience relation must be broadly logical supervenience if it is to be A-acceptable supervenience. The alternatives are to eliminate B-facts or else to give up A-ism (e.g., physicalism) concerning the B-facts. Neither C&J nor B&S countenance the elimination of facts about water, knowledge, or consciousness. Since elimination is assumed to be implausible, the locating facts must be logical or conceptual facts. Thus for C&J, B-facts must be conceptually or logically reducible to A-facts or else new sui generis connections are invoked—and the latter is unacceptable to physicalists.[21]

In the case at hand, the A-facts are facts of physics and we are concerned with physically acceptable supervenience—what Horgan (1993) calls *superdupervenience*. By the argument rehearsed above, the need for superdupervenience seems in turn to require reduction of the non-physics facts about water or consciousness to strictly physics facts using only broadly logical resources, just as C&J suppose. The alternative is to admit that the non-physics facts are additional facts, and thus to give up physicalism concerning both water and consciousness (see

[20] My discussion herein is also informed by the presentations and discussion at the conference, "Physicalism," at Bowling Green State University, Ohio, April 2005. See *Philosophical Studies* (October 2006, Vol. 131, Issue 1).
[21] In fact I think this argument equivocates on the class of "sui generis" facts. But that analysis goes beyond what I can offer herein. See my (forthcoming), in which I respond to Horgan's use of this basic argument.

Lynch and Glasgow, 2003.) For water, reduction is the only viable option because we all agree that water facts are physically acceptable facts outside of strict physics if there are any such.[22]

C&J's arguments thus concern the extent of the physically acceptable resources available to physicalists. Talk of A-facts, B-facts, and S-facts is just a way of keeping track of the "locations" of different facts and their relations. The important point is that the C&J model purports to tell us how to locate all physically acceptable facts. Hence, any facts that are commonly admitted as physicalistically acceptable but that don't fit C&J's model are counterexamples to their view. B&S argue that ordinary macrophysical facts, e.g., about water, are just such counterexamples. I concur. In the previous section I defended B&S, arguing that C&J have not established that facts about water are conceptually transparently reducible as they claim. This would make water facts sui generis according to C&J. So if B&S are right then consciousness is no different than water in this respect; neither is transparently reducible to physics alone. It seems that C&J have misunderstood the doctrine of physicalism. Specifically, it seems that they underestimate the resources available to physicalists by limiting them to the facts of physics along with some broadly logical facts.[23]

6. Armchair Knowledge

Above I asserted that the mere possibility that there is armchair knowledge that is not conceptual knowledge is enough to undermine C&J's argument. This is correct, for they have not provided a non-question-begging argument to rule out such knowledge. The best they can offer is that on

[22] Another response might be to welcome the anti-physicalist conclusion concerning water. But Chalmers cannot do this, for his arguments require that water differ from consciousness precisely in that it is reducible and consciousness is not.

[23] See also Polger (forthcoming).

their view we do not have to appeal to such knowledge.[24] But that conclusion itself depends on previously establishing their own view, and is in any case weaker than the required conclusion that physicalists *cannot* appeal to such knowledge.

This being said, I recognize that asserting that there could be armchair knowledge that is other than conceptual doesn't make it so. We will want to know more about such knowledge, or at least have some good *prima facie* examples. If armchair knowledge—of the sort on display when we make judgments about hypothetical cases involving water or knowledge—is not a priori and conceptual, then what sort of knowledge is it? Fortunately, we do have some ideas about what this armchair knowledge might consist in. In fact, the question of the kinds of and the justification of those beliefs or capacities that are drawn upon in the deployment of armchair "intuitions" or judgments has been a topic of lively philosophical discussion.[25] The issue is far reaching and vexed. Because my aim here is not to settle the matter but only to show that there is a matter to be settled, a few examples should be sufficient. I'll sketch four alternatives.

First, and perhaps most simply, we might think of armchair knowledge as tacit empirical knowledge or theorizing. This brings us back to B&S's critique of C&J, in which they point out that, "our reasoning about the proper epistemic response in various counterfactual situations is

---

[24] One might suppose that I carry the burden of showing not only that there is an alternative, but that there is an alternative that works as well as or better than C&J's favored transparent reductions at explaining, e.g., Gettier judgments. If C&J argued that their account is the best available, then that might be correct. But they do not—they argue that their account is the only candidate. In this context, then, I need only show that other alternatives are on offer. C&J might try responding to my alternatives by arguing that their account is superior on other grounds. But that would be a different argument than any that they have deployed thus far. (I am grateful to an anonymous referee for urging that I clarify this aspect of the dialectic.)

[25] Examples are too numerous to list in full. A few particularly salient sources not discussed in the text are: the exchange between Jackson and Harman in Michael and O' Leary Hawthorne (1994); the exchange between Boghossian and Harman in *Noûs* 30, 3 (1996); the exchange between Bealer, Lycan, and E. Sosa in *Philosophical Studies* 81 (1996); the essays collected in *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry* (DePaul and Ramsey, 1999); the symposium on Jackson's *From Metaphysics to Ethics* in *Philosophy and Phenomenological Research* 62, 3 (2001); the exchange between E. Sosa and Lynch in Greenough and Lynch (2006); as well as Bealer 1987, Lewis 1994, Fumerton 1999, Yablo 2000, Pust 2001, Kornblith 2006, D. Sosa 2006, E. Sosa 2007, Nolan (forthcoming).

informed not only by our concepts, but by implicit and explicit theories and general methodological principles that we have absorbed through our scientific culture—everything that the "we" who are performing these thought experiments believe"(43). The obvious implication is that, contra C&J, armchair judgments are not purely a priori. As I noted above, B&S settle for the weaker conclusion that whatever status such reasoning about water and knowledge have, a case can be made that consciousness is no different. Yet the stronger claim, that the judgments are not strictly a priori at all, deserves our attention.

A special case that of more or less tacit empirical theorizing is that of folk or commonsense understanding. Philosophers have had great interest in folk theories of psychology and physics, in particular. Such "theories" consist in beliefs about the subject matter that are regarded as self-evident by the holders. These beliefs may or may not be verbally articulable. Folk theories may contain genuinely conceptual knowledge, and even a priori knowledge—perhaps one example is the principle that a single object cannot be in two places at one time. But folk theories may also contain contingent beliefs about the way things are around here (e.g., that moss tends to grow on the north side of trees, or that a stitch in time saves nine), or that are demonstrably false.[26] Accuracy aside, it seems that folk knowledge about the world is a candidate for a extra-conceptual source of armchair judgments.

Of course, even granting that our judgments about hypothetical cases are empirically informed, C&J might hold that the empirical information takes the form of a conceptual constituent. And I have no doubt some tacit empirical knowledge is encoded in our concepts. But there is no reason to suppose that *all* common, commonsense, or folk knowledge will be encoded in the meanings of our concepts rather than simply recorded in ordinary tacit beliefs about the

[26] See McCloskey (1983) for nice examples of false folk beliefs about mechanics.

world. Perhaps on a very severe meaning-as-use theory one could hold that every judgment we make is constitutive of the meanings of the concepts employed in the judgment. But that view is problematic, and C&J do not defend such a doctrine in any event.

We have seen that one option is that armchair judgments involve tacit empirical knowledge or reasoning. As a second alternative, we could follow the suggestion of Horgan and Henderson (2002) mentioned above, by distinguishing between the knowledge that we have that is part of a concept and that which we acquire in the process of learning the concept. Horgan and Henderson think of both sorts of knowledge as conceptual; but for broadly Quinean reasons they don't think of either sort as a priori in the traditional "high grade" sense. As they explain the matter,

> Consider the typical situation in which one learns concepts. This is not a situation in which one's informants are interested solely, or even primarily, in one's acquiring the concept. Typically, the occasion for instruction arises when one is not capable of understanding some substantive point that the teacher wishes to convey. Perhaps one's instructor is trying to explain why a given monitor is better than another. The instructor mentions the pixel-density of the two, and gets a blank look in return. One apparently needs to acquire the concept of *pixel* in the course of this instruction. Instruction is then forthcoming both in certain concepts *and* in particular information about the monitors in question. And the two lessons may be given in a highly integrated and largely undifferentiated course of instruction. (2002)

While learning about computer monitors I acquire the concept *pixel*; and while acquiring the concept *pixel* and I learn about monitors. The knowledge about the monitors is not properly part

of the content of the concept *pixel*, but it may just as well be the basis for armchair judgments that appeal to no further empirical information. I may have the ability, given a description of a hypothetical situation, to judge that something is a computer monitor or that one monitor should be preferred over another. But the knowledge that fits this category may be plainly contingent. I may learn, while acquiring the concept pixel, that a company named 'Samsung' is one of the manufacturers of things that have pixels, for example. And even if we follow Horgan and Henderson in allowing that this knowledge is a kind of conceptual knowledge, it will not take the form of (or justify our assertion of) a priori knowable application conditionals. So this class of knowledge will not satisfying the requirements of transparent reductions. Yet C&J have not ruled out that our capacity to make judgments about water or knowledge is based on such concept-associated but not concept-constitutive knowledge.

The first two alternatives to C&J's model that I have considered propose that armchair judgments, for example those about Gettier cases, involve the expression of some tacit knowledge. This knowledge might have been acquired while acquiring concepts (Horgan and Henderson 2002) or it might simply be tacit empirical knowledge (B&S 1999). In both cases, however, the knowledge revealed in intuitions or armchair judgments can clearly be expressed in a propositional form. A third alternative is more open-ended on this last point. It may be that such judgments reveal not tacit beliefs but something less articulated, such as prototypes, biases, or associations. These might be conceptual or propositional, but they might also simply be aspects of our psychological mechanisms. They might be empirically acquired and even culturally elastic (Goldman 1989; Stich 1999; Weinberg, Nichols, and Stich 2001). This more psychological account has been part of the basis of the so-called experimental philosophy movement. Within the movement there is disagreement about whether philosophical judgments

reveal concepts, word meanings, beliefs, or whatnot (see, e.g., Alexander and Weinberg 2007, Knobe 2007, and Nadelhoffer and Nahmias 2007). Yet even if armchair judgments turn out to involve psychologized "conceptual" knowledge, such knowledge—like the liberalized "conceptual" knowledge admitted by Horgan and Henderson—would not be a priori in any strong sense and would not do the work that C&J require for their transparent reductions.

Finally, let us consider an entirely different approach: Timothy Williamson's (2004, 2005) hypothesis that the ability to make armchair judgments about hypothetical cases is a manifestation of a general cognitive capacity for dealing with counterfactuals. Williamson argues that the assessment of counterfactuals is an ordinary part of life rather than a special demand of philosophical reasoning: "Asserting counterfactuals is not distinctive of *a priori* methodology. They are often asserted on *a posteriori* grounds" (2005: 10). The cognitive ability to assess counterfactuals is similarly generic, according to Williamson:

> We have a general cognitive ability to handle counterfactual conditionals. When we have some conception of the circumstances in which *r* is true, and some conception of the circumstances in which *s* is true, we also have some conception of the circumstances in which the counterfactual $r \mathrel{\square\!\!\rightarrow} s$ is true. Sometimes we can reason from *r* to *s*, or from *r* to $\neg s$, perhaps using as auxiliary premises background beliefs that in some sense to be specified do not conflict with *r*. But for many cases 'reasoning' seems to imply a more formal or conceptually articulated process than we actually employ. Even concerning many counterexamples to philosophical analyses, our verdicts do not seem to be based on reasoning in any useful sense. Perhaps our ability to assess the truth-values of counterfactuals involves some capacity to simulate mentally the truth of the

> antecedent and to determine the truth-value of the consequent under that simulation, although just what that would involve is frustratingly obscure…. We can evaluate ['philosophical' counterfactuals] without leaving the armchair; we can also evaluate many 'unphilosophical' counterfactuals without leaving the armchair. After all, we had plenty of experience before we sat down. (Williamson, 2005: 13)

Williamson's proposal, then, does not explain our capacity to make judgments about hypothetical cases, e.g., concerning H_2O or knowledge, in terms of the mastery of concepts or the possession of application conditionals such as K or W. The counterfactual "reasoning" is not reconstructed as a valid deductive form with a tacit premise—be it a priori, conceptual, or otherwise. Instead Williamson holds that we have a cognitive capacity, perhaps supported by general knowledge of "how things go" in the world, that allows us to directly reason about counterfactual cases from the armchair. This proposal, then, provides a fourth alternative for how our reasoning about water, knowledge, and consciousness may be "informed not only by our concepts" and outstrip "just unfold[ing] our concepts" (B&S, 43).

Of course it could be asserted that the general capacity that Williamson appeals to is itself grounded in our mastery of certain a priori knowable application conditionals that are part of our conceptual capacities. Williamson himself seems to think not. For one thing, he denies that the capacity is a priori. And he also seems to doubt that the capacity to make counterfactual judgments is based on anything as articulated as the application conditionals, writing,

> Very often, the background knowledge needed to evaluate a counterfactual consists not of specific items of information acquired on specific occasions but of a more general sense of how things go, honed over long experience. Such a sense

> is typically not presented to the subject in usably verbal form; one says things like 'She would do that because she is that sort of person.' Of course, underlying the inarticulate sense of how things go must be some conformation of the brain, but the latter does not constitute a theory from which the subject can infer the counterfactual or its negation. (2005: 14)

Even if conceptual content need not be thought of as articulate, Williamson appears to resist the reduction of the counterfactual capacity to conceptual competence: "A few philosophers do deny that the subject in the Gettier case has justified true belief in the relevant sense without knowledge. Those philosophers exhibit theoretical deviance, perhaps bad epistemological judgement, but not linguistic incompetence. Some are native speakers of English; other native speakers of English do not classify them as incompetent at English. By any reasonable criterion, they understand the word 'know' and possess the concept *know*" (11-12). In Williamson's view, both reasoners posses the concept *knowledge*; and we may suppose that he takes the concepts to have the same contents.[27] So, pace C&J, it is not a difference in concepts that accounts for the disagreement between those who judge that Gettier scenarios are cases of knowledge and those who judge that they are not.

Needless to say, Williamson's hypothesized capacity to "handle" counterfactual conditionals demands further explanation. Williamson knows as much, commenting that "[s]urely we need a better epistemology of counterfactuals than we currently have" (13). And though there is even some empirical support for Williamson's idea (Byrne 2005), the present state of things leaves the nature of our ability obscure in a way that is surely incomplete. But that

[27] Of course for C&J the fact that the reasoners make different judgments constitutes their having different concepts, or rather their concepts having different intensions. But then the force of Williamson's proposal is as another demonstration that C&J cannot take for granted their preferred account of intensions.

mystery does not detract from my present appeal to Williamson's proposal. For one would have to *argue* contra Williamson that our capacity to assess counterfactuals after all depends on a priori conceptual knowledge of the sort glossed by application conditionals like W or K. That account could not be taken for granted as if there were no alternative but a conceptual explanation. Because C&J provide no such argument, Williamson's proposal stands as a prima facie alternative to C&J's model.

For my part, I am attracted to something like Williamson's proposal, so long as we do not suppose that our "capacity to handle counterfactual conditionals" is a special cognitive organ or module, perhaps located in the pineal gland. Like our "capacity to handle perceptual input" it is likely to employ quite general psychological mechanisms and information—no doubt even some concepts, at least on occasion. This idea is, in fact, not incompatible with the other three alternatives sketched above. Our capacity to deal with counterfactuals—whether about water, knowledge, or the weather foreast—likely engages empirical knowledge, folk theories, information acquired while learning concepts, conceptual knowledge, and innate and acquired prototypes and biases. In short, it probably engages all the kinds of cognitive and epistemic resources that human beings deploy to reason. This is, of course, a bit of empirical speculation. But it is such a generic and sketchy bit of speculation that one almost wonders how it could fail to be at least approximately correct.

This being said, my objection to C&J's reductio argument does not rely on any particular alternative account being established. No doubt objections can be raised against each of the proposals I have surveyed. But that fact only reinforces the central claim: C&J cannot assume that there is no available alternative to their proposal. The very fact that there is a dispute over the nature and justification of philosophical intuitions and judgments is evidence that C&J owe

us an argument that their way of interpreting our capacity for armchair reasoning about Gettier cases—and likewise about water and H_2O—is the only way. They are not entitled to suppose that their candidate runs unopposed. And, as I emphasized above, their dialectical obligations will not be met simply by arguing that their view could do the work that its competitors can do. They must establish the stronger claim that physicalists cannot consistently appeal to any other account. This they have not done.

7. Taking Stock

I have shown that C&J's rebuttal to B&S depends on a reductio argument that compares what can be "deduced" about water from facts about H_2O to what can be "deduced" about knowledge from information about various hypothetical situations. I suggested, developing ideas suggested by B&S and by Lycan, that such knowledge (even if it were a priori) may not be strictly conceptual, and may not be transparently "entailed" by the base facts alone. If so, then the fact that facts about consciousness cannot be "deduced" from physical facts does not distinguish them from any other macroscopic or philosophical facts, such as facts about water or knowledge. This is a serious setback for Chalmers' anti-physicalist arguments. And it is a setback for Jackson's argument that physicalism can only be vindicated by conceptual analysis.

If my presentation of the core commitments of C&J is correct, then we can now see that one of the reasons that their arguments are both highly polarizing and difficult to assess is that they are forcing critics into a metaphilosophical debate about the nature and content of philosophical knowledge. Despite appearances, their arguments do not depend only on a first-order question concerning the logical structure of explanation, but also on a higher-order question about what resources are available for philosophical reasoning. Taken in this context,

the failure of C&J's reductio argument may undermine their general campaign for the centrality of a kind of conceptual analysis for all philosophical reasoning.

I have made a first pass at showing that physicalists could have more resources than C&J suppose.[28] I concede that more must be done to positively establish the nature and availability of extra-conceptual resources for physicalist metaphysics. But for now I will satisfy myself with the conclusion that C&J have not ruled out all alternatives to their own view.

8. References

Alexander, J. and J. Weinberg. 2007. Analytic Epistemology and Experimental Philosophy. *Philosophy Compass*, 2 (1): 56-80.

Bealer, G. 1987. The Philosophical Limits of Scientific Essentialism. *Philosophical Perspectives*, 1: 289-365.

Bealer, G. 1996. A Priori Knowledge and the Scope of Philosophy. *Philosophical Studies*, 81: 121-142.

Bealer, G. 1996. A Priori Knowledge: Replies to William Lycan and Ernest Sosa. *Philosophical Studies*, 81: 1163-174.

Block, N. and R. Stalnaker. 1999. Conceptual analysis, dualism, and the explanatory gap. The *Philosophical Review*, 108 (1): 1-46.

Boghossian, P. 1996. Analyticity Reconsidered. *Noûs*, 30 (3): 360-391.

Byrne, A. 1999. Cosmic Hermeutics. *Philosophical Perspectives* 13: 347-383.

Chalmers, D. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

[28] This project is further pursued in my forthcoming.

Chalmers, D. and F. Jackson. 2001. Conceptual analysis and reductive explanation. The *Philosophical Review*, 110 (3): 315-361.

Crane, T. and H. Mellor. 1990. There is No Question of Physicalism. *Mind*, 99:185–206.

Dowell, J. unpublished manuscript. Serious Metaphysics and the Vindication of Explanatory Reduction.

Fumerton, R. 1999. A Priori Philosophy after an A Posteriori Turn. *Midwest Studies in Philosophy: XXIII*: 21-33.

Gettier, E. 1963. Is justified true belief knowledge? *Analysis* 23: 121-123.

Goldman, A. 1989. Psychology and Philosophical Analysis. *Proceedings of the Aristotelian Society* 89:195–209.

Greenough, P. and M. Lynch (eds). 2006. *Truth and Realism*. Oxford: Oxford University Press.

Harman, Gilbert. 1994. Doubts About Conceptual Analysis. In Michaelis Michael and John O' Leary Hawthorne (eds), *Philosophy in Mind* (Kluwer): 43-48.

Harman, G. 1996. Analyticity Regained? *Noûs*, 30 (3): 392-400.

Hawthorne, J. 2002. Advice for Physicalists. *Philosophical Studies* 108: 17-52.

Hellman, G. and F. Thompson.1975. Physicalism: Ontology, Determination, and Reduction. *Journal of Philosophy*, 72: 551–564.

Hill, C. 1998. Chalmers on the A Priority of Modal Knowledge. *Analysis* 58 (1): 20-26.

Hill, C. and B. McLaughlin. 1999. There Are Fewer Things in Reality than Are Dreamt of in Chalmers' Philosophy. *Philosophy and Phenomenological Research* 59: 445-54.

Horgan, T. 1984. Supervenience and Cosmic Hermeneutics. *Southern Journal of Philosophy*. Supp. 22:19-38.

Horgan, T. 1993. From Supervenience to Superdupervenience: Meeting the Demands of a Material World. *Mind* 102: 555-586.

Horgan, T. and D. Henderson. 2002. The a priori isn't all that it's cracked up to be, but it is something. *Philosophical Topics* 29: 219-50.

Jackson, F. 1982. Epiphenomenal qualia. *Philosophical Quarterly,* 32:127-36.

Jackson, Frank. 1994. Armchair Metaphysics. In Michaelis Michael and John O' Leary Hawthorne (eds), *Philosophy in Mind* (Kluwer): 23-42.

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

Jackson, F. 2005. The case for A Priori Physicalism. Christian Nimtz and Ansgar Beckermann (eds.) (2005). Philosophy – Science – Scientific Philosophy. Main Lectures and Colloquia of GAP.5, Fifth International Congress of the Society for Analytical Philosophy, Bielefeld, 22–26 September 2003, Paderborn: mentis.

Knobe, J. 2007. Experimental Philosophy. *Philosophy Compass*, 2 (1): 81-92.

Kornblith, H. 2006. Appeals to Intuition and the Ambitions of Epistemology. In Stephen Hetherington (ed), *Epistemology Futures* (London: Oxford University Press): 10-25.

Levine, J. 1983. Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly*, 64, 354-361.

Lewis, D. 1994. Lewis, David: Reduction of Mind. In Samuel Guttenplan (ed), *A Companion to the Philosophy of Mind*. Oxford, UK: Blackwell Publishers: 412-431.

Lycan, W. 1996. Bealer on the Possibility of Philosophical Knowledge. *Philosophical Studies*, 81: 143-150.

Lycan, W. (forthcoming). Serious Metaphysics: Frank Jackson's Defense of Conceptual Analysis. In Ian Ravenscroft (ed.), *Minds, Worlds and Conditionals: Essays in Honour of Frank Jackson* (Oxford: Oxford University Press).

Lycan, W. 2003. Vs. a New A Priorist Argument for Dualism. In E. Sosa and E. Villanueva (eds.), *Philosophical Issues*, Vol. 13: 130-47.

Lynch, M. and J. Glasgow. 2003. The Impossibility of Superdupervenience. *Philosophical Studies*, 113: 201-223,

McCloskey, M. 1983. Naive Theories of Motion. In Gentner, D. and Stevens, A. (eds.), *Mental Models*. Hillsdale, NJ: Erlbaum.

McLaughlin, B. 2005. A Priori Versus A Posteriori Physicalism. Christian Nimtz and Ansgar Beckermann (eds.) (2005). Philosophy – Science – Scientific Philosophy. Main Lectures and Colloquia of GAP.5, Fifth International Congress of the Society for Analytical Philosophy, Bielefeld, 22–26 September 2003, Paderborn: mentis.

Nadelhoffer, T. and E. Nahmias. 2007. The Past and Future of Experimental Philosophy. *Philosophical Explorations*, 10:2, 123-149.

Nolan, D. forthcoming. Platitudes and Metaphysics. In David Braddon-Mitchell & Robert Nola (eds), *Naturalism and Analysis* (MIT Press).

Polger, T. forthcoming. Moorean Supervenience.

Pust, J. 2001. Against Explanationist Skepticism Regarding Philosophical Intuitions. *Philosophical Studies*106: 227-258.

Sosa, D. 2006. Scepticism About Intuition. *Philosophy* 81: 633-647.

Sosa, E. 1996. Rational Intuition: Bealer on its Nature and Epistemic Status. *Philosophical Studies*, 81: 151-162.

Sosa, E. 2007. Experimental Philosophy and Philosophical Intuition. *Philosophical Studies*132: 99-107.

Stich, S. 1999. Reflective Equilibrium, Analytic Epistemology and the Problem of Cognitive Diversity. In Michael DePaul and William Ramsey (eds.), *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry* (Lanham, MD: Rowman and Littlefield): 95-112.

Weinberg, J., Nichols, S. and Stich, S. 2001. Normativity and Epistemic Intuitions. *Philosophical Topics*, 29: 429 - 460.

Williamson, T. 2004. Philosophical 'Intuitions' and Scepticism about Judgement. *Dialectica* 58 (1): 109-153.

Williamson, T. 2005. Armchair philosophy, metaphysical modality and counterfactual thinking. *Proceedings of the Aristotelian Society* 105, 1: 1-23.

Wilson, J. 2005. Supervenience-based Formulations of Physicalism. *Nous* 39 (3): 426-459.

Wright, W. 2007. Explanation and the Hard Problem. *Philosophical Studies* 132: 301-330.

Yablo, S. 2000. Textbook Kripkeanism and the Open Texture of Concepts. *Pacific Philosophical Quarterly* 81: 98-122.

Yablo, S. 1999. Concepts and Consciousness. *Philosophy and Phenomenological Research* 59: 455-63.