

# Empirical Metaphysics: The Role of Intuitions about Possible Cases in Philosophy<sup>1</sup>

J. L. Dowell

Draft at 12/11/07

Here's an interesting fact about speakers on which many arguments for and against a variety of philosophical positions depend: We often find ourselves with strong opinions about the proper evaluation of conditionals constructed in accordance with the following recipe.

Step 1: Identify an expression for the object of some investigation, e.g. 'physical', 'knowledge', 'right', 'free will', 'cause'...the list goes on.

Step 2: Letting "E" be the expression in question, describe a possibility,  $W_n$ , in such a way that it is not settled explicitly by that description which things are E in  $W_n$ .

Step 3: Suppose that the actual world,  $W_a$ , is  $W_n$ , and ask oneself what one wants to say that E is in  $W_n$ .<sup>2</sup>

In other words, see whether one can fill out the "?" in a conditional with the following form in such a way that makes the whole conditional one that one prepared to accept:

(C1) If  $W_a=W_n$ , then, if anything,  $E=?$ <sup>3</sup>

Call conditionals that have this form, *intuition-pumping conditionals* or *intuition-pumpers* for short. According to Frank Jackson, intuition-pumping conditionals provide a quasi-formalization of the common philosophical practice of asking individuals to characterize scenarios designed to be of philosophical interest.<sup>4</sup> He suggests,<sup>5</sup> for

---

<sup>1</sup> Many thanks to David Braddon-Mitchell, Al Casullo, Sean Foran, Gene Witmer, Jonathan Schaffer, Joshua Schechter, David Sobel, Ryan Wasserman, and Jessica Wilson for helpful discussions of the issues raised in this paper. Thanks also to the participants in the discussion of earlier versions of this paper at the University of Nebraska-Lincoln and at the BSPC 2007.

<sup>2</sup> Thanks to John Gibbons for discussion here.

<sup>3</sup> Where appropriate, "E=?" is replaced by a biconditional formula, e.g. where "E" is replaced by a predicate.

<sup>4</sup> See Jackson [1998a] pp.28-37 and 47-8.

<sup>5</sup> See, for example, his [1998a] pp.28-55 and [1998b].

example, that they are central to one of Putnam's twin earth thought experiments. (Question: If we're on earth and the stuff in lakes and oceans on earth is H<sub>2</sub>O and the stuff in the lakes and oceans of Twin Earth (a planet remotely located in the actual world) is XYZ, water on Twin Earth=? Answer: H<sub>2</sub>O.) Similarly, Jackson suggests,<sup>6</sup> in Gettier's arguments, the evaluation of such conditionals convinced most people that justification, truth, and belief were by themselves insufficient for knowledge.

It is easy to add to Jackson's examples. In the case of the trolley problem, we want to know whether it's ever morally permissible to knowingly take a human life.<sup>7</sup> In the responsibility debates, cases involving counterfactual interveners are to help us understand whether responsibility requires that one could have done otherwise.<sup>8</sup> Wesley Salmon's example of a shadow-dot making a circuit around a circular room is to show that counterfactual analyses of causation are insufficient.<sup>9,10</sup> All of this is but a reminder

---

<sup>6</sup> Jackson [1998a] pp.28-9.

<sup>7</sup> See Thomson [1985] p. 84.

<sup>8</sup> Frankfurt [1969].

<sup>9</sup> Salmon [1993].

<sup>10</sup> Jackson's claims here are arguably oversimplified. One might dispute his claim that, in the discussions of particular intuitions just mentioned, anything quite as specific as the evaluation of a Jackson-style intuition-pumper is in play. This is unfortunate, I think, not for Jackson's argument, but for the reliance on intuitions in those discussions. Care requires that we distinguish three different sorts of intuitions, on the one hand, and three different sorts of account that such intuitions might be thought to support or undermine, on the other.

Consider three different sorts of account. The form of an account of the kind at issue in many philosophical discussions is 'Fs are Gs'. Accounts of this very generic kind differ as to whether they are 1) contingent, a posteriori universal generalizations, 2) a priori, or 3) metaphysically necessary. Post-*Naming and Necessity*, most philosophers regard a priority and metaphysical necessity as capable of coming apart. Given this, we should expect an in principle difference in the kinds of scenario or possibility that undermine each of these accounts. Unfortunately, the philosophical practice of appealing to intuitions to become entrenched long before Kripke's book and has altered little since.

So, to be more clear: We need actual counter-cases to undermine universal generalizations and we need intuitions about what is epistemically possible (the kind that Jackson's intuition-pumpers are to help us discover) to undermine putative a priori truths. And, of course, we need intuitions about what is possible counterfactually to undermine purported metaphysical necessities.

So, which are at issue in the sample philosophical disputes I've mentioned? The papers I've cited don't mark the distinctions that I have just marked. Nonetheless, I think there are two sources of evidence for the view that something like epistemic possibility is in play in the intuitions their authors appeal to. First, all of the scenarios these authors imagine are described by them in the present tense, rather than the counterfactual-indicating subjunctive. (E.g. Salmon does not say 'suppose that there *had been* a spotlight

that it is a philosopher's default assumption that widespread agreement in our intuitions counts as evidence for theories that accommodate them. Indeed, the practice is so widespread as to deserve being thought of as part of a philosopher's stock-in-trade.

The question here is: How must we understand our evaluation of intuition-pumping conditionals such that they are able to play the role of providing evidence for and against philosophical positions? One recent answer to this question has been defended by Frank Jackson and David Chalmers. According to Jackson and Chalmers, the best explanation of our ability to evaluate such conditionals is that competent speakers have implicit a priori knowledge of what fixes the extensions of many of their terms. Jackson and Chalmers each separately defend the further claim that these extension-fixers are components of the overall contents of such terms and offer detailed proposals for how best to understand the nature of such contents. Such a priori extension-fixers, in turn, provide substantive answers to philosophical questions about the nature of the physical, the right, knowledge, etc. by telling us what something has to be like in order to fall within the extension of our terms for these targets of our investigations.

By vindicating our reliance on intuitions in settling philosophical disputes, this explanation of our truth-sorting capacities is a powerful source of evidence in favor of the Jackson/Chalmers position. Indeed, this explanation creates a burden for those who would reject that position:<sup>11</sup> Either offer an equally plausible explanation that doesn't require a priori knowledge of extension-fixers or give up reliance on intuitions. The

---

mounted in the center of a circular room...'. Rather, he says "suppose that *we have* a very large circular building...with a spotlight mounted at its centre" (Salmon [1993] p. 153. See also Gettier [1963] p.122.) Second, these authors often describe the theories they aim to undermine as purported analyses or a priori truths. (See for example Frankfurt [1969] p.829,)

Rather than quoting extensive passages from Gettier, Frankfurt, Salmon, and Thompson, I invite the reader to check these characterizations for herself.

<sup>11</sup> Foes of a priori extension-fixing components of content include Byrne and Pryor [2006], Bigelow and Schroeter [forthcoming], Block and Stalnaker [1999], Soames [2005], Stalnaker [2003] and Yablo [2000].

entrenched reliance on intuitions in philosophical practice makes the first of these disjuncts the more promising tack.

The central aim here is to do just that, to find a way to vindicate our reliance on intuitions without requiring a priori extension-fixers. My strategy here will be to grant for the sake of argument three assumptions on which their arguments depend. First, I'll assume that our referential intentions or policies have an important role to play in determining the extensions of many of our terms. This assumption isn't terribly controversial. Indeed, it's neutral between a variety of positions on reference-fixing.<sup>12</sup> Second, Jackson and Chalmers must both assume that there is a tight connection between our referential intentions and our disposition to apply our expressions. (This 'tight connection' might be thought of in a variety of ways. One might hold that the disposition just is our implicit intention or that the intention guides the disposition or that the disposition is a disposition to be guided by the intention.) Third, both must assume our descriptive reactions to scenarios considered as actual provide our best evidence for the contents of those extension-determining dispositions.

I will argue that if we grant these central assumptions, we should regard our evaluations of intuition-pumping conditionals as providing an important, a priori source of evidence for hypotheses about the contents of our implicit referential policy with regard to a term. But this a priori source of evidence is empirically defeasible. In other words, gold-standard evidence for such hypotheses is empirical. Whether this means that our knowledge of our reference-fixers is always a posteriori depends in part on what we want to say about at least one of two issues. The first is the question of what sort of

---

<sup>12</sup> It's neutral, for example, between Kripke's, Putnam's, and Stalnaker's positions as well as Jackson's and Chalmers'.

justification one thinks knowledge requires. One who holds that genuine knowledge requires gold-standard evidence will say that such knowledge is always a posteriori. One who holds that such evidence isn't always required may hold such knowledge can be a priori.

The second issue is whether to hold that properly a priori knowledge could have an empirical falsifier. This issue is at least partly a terminological one about how best to categorize kinds of knowledge. The issue is a large one and settling it (or the first issue) is well beyond the scope of this paper. However, in section III, I'll briefly give a few reasons not to call our knowledge of our referential policies "a priori".

This array of plausible positions on empirical defeaters does mean that we *needn't* posit implicit a priori knowledge of reference-fixers in order to explain our confident evaluations of intuition-pumpers or to vindicate our reliance on those evaluations in philosophical theorizing. In contrast to the Jackson/Chalmers proposal, I'll argue that if we carefully observe what we're doing in the evaluation of intuition-pumping conditionals, we'll see that they are better understood as providing reliable evidence for the development of plausible armchair hypotheses about the contents of our referential dispositions and policies, hypotheses that await empirical evidence for their ideal confirmation.

The differences between the Jackson/Chalmers position and the one defended here may seem minor. But in fact, if the position defended here is correct, their central case for a priori knowable components of content fails and so too their case for the claim that conceptual analysis is required for the defense of reductions.<sup>13</sup>

---

<sup>13</sup> See their [2001].

## I. The Jackson/Chalmers Proposal

In their details, Jackson's and Chalmers's proposals differ as to how best to understand what these a priori components of content are, exactly, and how those contents are related to the much-discussed two-dimensional semantic framework. But at the core of each of their proposals is the same set of claims. They are:

1. Competent users of a term have implicit a priori knowledge of the extension-fixers for many of their terms.
2. These extension-fixers are an a priori known kind of content (i.e. primary or A-intensions).
3. Competent speakers' confident evaluations of intuition-pumping conditionals are best explained by positing implicit, a priori knowledge of those extension-fixers and thus, a kind of the content (by #2).
4. Given #3, competent speakers' evaluations of such conditionals provide a central source of evidence for the truth of #1.
5. Using our evaluations of such conditionals we are in principle able to discover interesting a priori truths by finding two distinct expressions that share an extension-fixer and constructing an identity statement between the two, e.g. that water =the actual, local wet stuff that is predominantly found in the lakes and oceans.

For ease of exposition, I'll call 1-5 "the Jackson/Chalmers proposal" or "the Jackson/Chalmers programme", but it should be kept in mind that there are significant

differences between their proposals, albeit ones that do not touch upon the issues raised here.<sup>14</sup>

What is their evidence for these claims? Here too there are differences in detail between each of their arguments that can be glossed over in favor of a more general statement. Perhaps the most compelling statement of that general argument is the following:

Let “E” be some expression and S, any speaker that is only minimally semantically competent with “E”. Also, let “E”’s extension-fixer be whatever (possibly disjunctive) property or set of properties something has to have in order to be in “E”’s extension. (However highly unnatural the set that makes up “E”’s extension, there will be some such disjunctive property, in a relaxed Lewisian sense.) Finally, let a semantically neutral characterization  $C_n$  of a world  $W_n$  be a description of  $W_n$  that does not explicitly specify what the semantic truths<sup>15</sup> or E-truths (truths specified using “E”) are at  $W_n$ . (This condition is required so that S’s evaluations of intuition-pumpers can provide data for semantic hypotheses. If the antecedents explicitly settled the truth of the hypotheses that are at issue, then S’s evaluations couldn’t be evidence for or against those hypotheses.)

1. If S could in principle recognize the extension of “E” in any world  $W_n$  supposed to be actual, given a semantically neutral characterization  $C_n$  of  $W_n$ , then S implicitly knows the extension-fixer for “E”<sup>16</sup> for any world

---

<sup>14</sup> There is one difference between their views that matters for what they will each need to say for a certain range of terms. I discuss this difference in footnote 27. This difference is worth discussing, but does not matter for the purposes of formulating my central case against the shared core of their positions.

<sup>15</sup> These are truths that about what the extension of “E” is in  $W_n$ .

<sup>16</sup> Jackson’s position is that our terms have a priori knowable components of content, while Chalmers’s position is that our term-tokens do. For ease of exposition, I put the argument in Jackson’s terms. But the argument, put in terms of S’s tokenings of “E”, fits Chalmers’s position as well.

supposed to be actual in which “E” has a determinate extension.

(Henceforth this last caveat should be taken as read.)

2. If S knows the extension-fixer for “E” for any world supposed to be actual, then there is an a priori knowable component of “E”’s content.
3. There are many expression-types such that, for any world  $W_n$  supposed to be actual, we can in principle recognize the extension of each expression of those types from a semantically neutral description  $C_n$  of  $W_n$ .
4. Therefore, many of our expressions have an a priori knowable component of their content.<sup>17</sup>

The argument is clearly valid. So if the conclusion is false, at least one of the premises must be false as well. Which one? The second premise must be true—it gives a sufficient condition for a term’s having an a priori, implicitly known kind of content in the Jackson/Chalmers sense, i.e. a kind that is known from minimal semantic competence alone. Here I’ve substituted “supposed to be actual” for their locution “considered as actual” to make it a bit more clear how, on their view, armchair experiments are to yield a priori knowledge. On a natural understanding, one way of considering a world as actual is believing that it is. But a priori knowledge that P in their sense requires that P is knowable while suspending all belief as to which world is actual. Their idea is that whatever is knowable regardless of which world is actual is knowable without knowing anything about the empirical features that distinguish the actual world from other merely possible ones. And what is knowable without knowing anything about the empirical features that distinguish the actual world from the other merely possible ones is a priori knowable in one plausible sense. So if S knows E’s extension-fixer for any world

---

<sup>17</sup> See Jackson [1998] and Chalmers [2004].

supposed to be actual, she knows E's extension-fixer a priori, that is, she knows what property or set of properties something has to have in order to be in E's extension, regardless of which world turns out to be the one she is in.

The first premise might look like the weak one. It might seem that a speaker could be minimally semantically competent with "E", able to recognize "E"'s extension in any world supposed to be actual, and yet fail to know, even implicitly, what it is that fixes "E"'s extension. But given what Jackson and Chalmers mean by "implicit knowledge", the first premise must be true. For them, implicit knowledge confers a recognitional capacity that one needn't be able to give an account of. For example, implicit knowledge of what it takes to be a well-formed formula in propositional logic would confer on one the ability to distinguish well- from ill-formed formulas, even if not the ability to explicitly formulate the difference between the two. So, if S doesn't know what the extension-fixer is for "E", she doesn't know, even implicitly, what property or set of properties something has to have in order to be in "E"'s extension. Let S now be given a semantically neutral characterization  $C_n$  of some world  $W_n$ .  $C_n$  will give a complete description of  $W_n$ , of everything that might serve as a potential extension for "E" in  $W_n$ . Given that  $C_n$  is semantically neutral, without at least implicit knowledge of what fixes "E"'s extension, S has no way of identifying "E"'s extension in  $W_n$ . Since this is true for any world we take, it follows that if the consequent of the first premise is false, so is its antecedent; in other words, it follows that the conditional itself is true.

I'll argue that the weak premise is the third one. While we may be able to formulate highly reliable hypotheses about what the extensions for some of our expressions are while suspending belief as to which world is actual, we not in a position

to know what those extensions are in any way that implies that we thereby have implicit, a priori knowledge of the extension-fixers for those expressions. (For ease of exposition, this is a quick version of my conclusion. I'll put the point a bit more carefully below.)

## **II. The Jackson/Chalmers Argument for A priori Extension-fixers**

What might be said in defense of the third premise? Consider again the form of intuition-pumping conditionals:

(C1) If  $W_a=W_n$ , then, if anything,  $E=?$ <sup>18</sup>

On the Jackson/Chalmers view, when evaluating such conditionals, we treat ourselves as instruments for collecting data that support hypotheses about what the content of our implicitly adopted referential policy for some expression “E” might be. As input, we suppose that it turns out that the actual world is qualitatively just like some world  $W_n$  according to its semantically neutral characterization  $C_n$ . Then we measure our willingness to recognize items as E in  $W_n$ .<sup>19</sup> The output, our descriptive reaction, counts as evidence for hypotheses about the contents of our referential policies because they are produced by our descriptive dispositions triggered by the input. (For brevity's sake, I'll sometimes call such hypotheses “referential hypotheses”.) Find the pattern in a speaker's descriptive reactions and you will find what something has to be like in order for it to fall within the extension of “E” for that speaker, given her implicitly held referential policy governing her use of “E”.

---

<sup>18</sup>See footnote 3.

<sup>19</sup> Chalmers [2004] pp.176-184. See also Jackson [1998] p. 36: “The business of extracting the cases that count as Ks [where Ks are the target of some investigation] from a person's responses to possible cases is an exercise in hypothetico-deduction. We are seeking the hypothesis that best makes sense of their responses taking into account all the evidence.”

For some filling-in N of “E=?”<sup>20</sup> to make explicit S’s implicitly held referential policy for “E” two things must be true. First, there must be no  $W_n$  such that ‘If  $W_a=W_n$ , then  $E=N$ ’ makes  $E=N$ , and so the whole conditional, come out false. Second, in evaluating intuition-pumpers, our speaker must suspend all belief as to which world is actual. Since to be a priori in the Jackson/Chalmers sense is to be knowable while suspending all belief as to which world is actual, the input our measuring device is fed must be worlds supposed, but not believed, to be actual.<sup>21</sup>

To illustrate how their procedure is supposed to work, consider a very simple example. Suppose we are trying to find out what S’s referential policy is for “water” and so what property or set of properties something has to have in order to count as water for S. Following their procedure, we ask S to consider a set of conditionals such as:

C2: If  $W_a$ =the world in which the actual wet stuff in the local<sup>22</sup> lakes and oceans is  $H_2O$  and..., then, if anything, water in  $W_a$  =?

C3: If  $W_a$ =the world in which the actual wet stuff in the local lakes and oceans is XYZ and..., then, if anything, water in  $W_a$  =?

Suppose after asking S to evaluate a variety of such conditionals, we notice a pattern emerging. We notice that for any  $W_n$ , whatever the wet stuff in the local lakes and oceans in  $W_n$  turns out to be S recognizes upon careful reflection as water in  $W_n$ . We then ask S to evaluate the following further intuition-pumpers:

---

<sup>20</sup> See footnote 3.

<sup>21</sup> Chalmers [2004], p.177. “On the epistemic understanding, to consider a world as actual is to consider the hypothesis that W is one’s own world.” See also, Jackson [1998] p. 51: ‘What we can know independently of knowing what the actual world is like can properly be called a priori.’

<sup>22</sup> In this case, to evaluate the consequent on the supposition that the antecedent is true requires that that antecedent include information about where S is located in  $W_n$ .

C4: If  $W_a$ =the world in which the actual wet stuff in the local lakes and oceans is  $H_2O$  and..., then, if anything, water in  $W_a$  = the wet stuff in the local lakes and oceans.

C5: If  $W_a$ =the world in which the actual wet stuff in the local lakes and oceans is XYZ and..., then, if anything, water in  $W_a$  =the wet stuff in the local lakes and oceans.

Suppose further that after considering a number and variety of such conditionals, we notice that whenever a world contains wet stuff in its local lakes and oceans, S calls whatever that stuff is “water”. In other words, from the armchair, S treats “water” and “the wet stuff in the local lakes and oceans” as having the same extension in every world supposed to be actual. What do the results of this experiment show?

For Jackson and Chalmers, such data show, first, that S is able to recognize the extension of “water” for any world supposed to be actual. Second, they show that S has adopted the same referential policy for “water” as she has for “the wet stuff in the local lakes and oceans”. Third, since S is able to establish this result from the armchair consideration of intuition-pumpers alone, this shows that she has implicit a priori knowledge of the contents of her referential policies for each and so that she has implicit, a priori knowledge that water = the actual wet stuff in the local lakes and oceans.

More generally, such experiments are to show that minimally competent speakers have implicit, a priori knowledge of: 1) an extension-fixing component of content for many of their terms and 2) certain truths, expressible using terms for which they have adopted the same extension-fixing policy.<sup>23</sup> Their implicit, a priori knowledge can be made explicit by finding the pattern among their armchair descriptive reactions that

---

<sup>23</sup> See, for example, Jackson [1998a], [1998b], and [2001] and Chalmers [2002b] and [2004].

reveals their descriptive disposition. Speakers can decide between rival referential hypotheses by seeing which one most plausibly guides their descriptive disposition. Since all of this reasoning can be done from the armchair, if correct, this story shows how we can have a priori knowledge of our reference-fixing policies *and* does so in a way that vindicates our reliance on intuitions in philosophical theorizing. In other words, if Jackson and Chalmers are right about what explains our ability to evaluate intuition-pumping conditionals, then the argument above for a priori accessible components of content is not only valid, but sound.

### **III. Intuition-pumping Conditionals and the Evidential Weight of Their Evaluation**

In fact, though, the argument is not sound, or, if it is, its soundness isn't established by the considerations as yet marshaled by Jackson and Chalmers in its defense. I'll argue that if we really think about what our armchair ability to evaluate intuition-pumpers comes to, we'll see that that ability need not be understood as a manifestation of our a priori knowledge of an extension-fixing kind of content. That's because when our actual descriptive reaction on coming to believe that a certain scenario is actual comes apart from our armchair reaction while supposing that scenario to be actual (and a few other conditions discussed below are met), we should treat our actual reaction as undermining the evidential status of our armchair reaction. I'll run through the basic argument for these claims in this section and then consider a series of objections and replies to them in the next.

Let me first emphasize that my aim in this section is to lay out a rival explanation of the source of the evidential weight of our armchair intuitions, *not* to prove that the Jackson/Chalmers explanation is incorrect. The position outlined and defended here is a

rival to their explanation that does not rely on positing the existence of speakers' a priori knowledge of extension-fixing contents. My primary aim is to defend the plausibility of that rival explanation. In my view, the considerations marshaled in its defense (both in this section and in section IV) make it more plausible than the Jackson/Chalmers explanation. But my primary conclusion does not require anything so strong. My primary conclusion is in this respect weak: I claim not to have proved that the Jackson/Chalmers explanation is incorrect. Rather, my primary conclusion is that the plausibility of the explanation I outline suffices to reveal an important gap in the Jackson/Chalmers argument for a priori knowledge of extension-fixers. The Jackson/Chalmers claim is that we *need* such extension-fixers to vindicate the evidential value of our armchair intuitions. The plausibility of the present position shows that we do not need any such thing.

Back to business. Notice first that in order for the Jackson/Chalmers proposal to be at all plausible, it will need to be compatible with the possibility of mistakes in our armchair acceptance of referential hypotheses. And their proposal is so compatible. Since according to that proposal our referential policies are implicit, it is possible to be mistaken about their contents. Neither Chalmers nor Jackson supposes that immediate introspection will result in a clear verdict on their contents, e.g. the content of one's referential policy with respect to the term "water".<sup>24</sup> On their view, in order for truth-evaluations of intuition-pumpers to help reveal a speaker's referential policy, those evaluations must be made by that speaker while 'ideally rational'. Here are two possible sources of error, stemming from irrationality, that do not undermine the

---

<sup>24</sup> See, for example, Chalmers' [2006] p. 91 and his footnote 13 reply to an objection by posed by Byrne and Pryor [2006].

Jackson/Chalmers project: First, if our evaluations of a particular series of intuition-pumping conditionals are hasty, we may judge their consequents true or false at a world supposed to be actual when on more careful reflection we would not. By Jackson's and Chalmers' lights, such false 'data' could lead to mistakes in our identification of the pattern that captures our true referential disposition. Second, similar mistakes may result if the number of worlds supposed to be actual is too low. In such a case, we may infer that our usage is guided by a simpler pattern in our expression's extension than it really is.

The first lesson, then, is that not all of our armchair descriptive reactions provide evidence for hypotheses about our extension-fixing intentions. Hasty reactions, for example, don't. The second lesson is that, even with prima facie good armchair evidence, we can be led astray in our hypothesis selection. Failure to consider a sufficiently large number and variety of intuition-pumpers, for example, can do that. Importantly, both of these sources of error are discoverable from the armchair. That's why they don't threaten the Jackson/Chalmers project. Are other sorts of mistakes possible? In particular, are there possible errors not discoverable from the armchair? I'll argue that one source of empirical defeasibility not stemming from irrationality is inherent in the armchair evaluation of intuition-pumping conditionals.

To help focus the issue, think about a very natural way of understanding what is involved in supposing that a certain world,  $W_n$ , is actual. In supposing that  $W_a=W_n$ , we seem to be pretending to believe that  $W_a=W_n$ . What kind of attitude is pretend belief? On this natural understanding, when we pretend to believe that P, we're in a state that aims to mimic the belief that P. (Some evidence for this natural understanding is that

when we suppose that P, we allow P to figure in our reasoning in a way similar to the way a belief that P would.<sup>25</sup>)

Such a pretense may meet its aim with varying degrees of success. The more closely a pretend-belief that P mimics a genuine belief that P, the more it succeeds as a pretend-belief. So, in supposing that P, we seem to be trying to think ourselves into a state that mimics the belief that P. The point of pretending to believe that P, we might say, is to get us into the shoes of someone who believes that P. This suggests an explanation for why our armchair reactions can provide evidence for hypotheses about our descriptive dispositions: Our armchair descriptive reactions can provide evidence for such hypotheses precisely because it is plausible to think that if we pretend to believe that  $W_n$  is actual well enough, we'll trigger our referential dispositions by creating conditions that sufficiently mimic how we'd describe  $W_n$  on coming to believe that  $W_n$  is actual.

On this understanding of our armchair evaluations of intuition-pumpers, then:

1. Supposing that  $W_a=W_n$  is a kind of pretend belief that  $W_a=W_n$ .
2. S's pretending to believe that  $W_a=W_n$  at time  $t_1$  aims to mimic (at least in the relevant respects) the state that S would be in were S to believe that  $W_a=W_n$  at  $t_1$ .<sup>26</sup>
3. A pretend-belief<sub>1</sub> that  $W_a=W_n$  (PB<sub>1</sub>) is a more successful pretend-belief that  $W_a=W_n$  than pretend-belief<sub>2</sub> (PB<sub>2</sub>) if PB<sub>1</sub> more closely resembles in the relevant respects a genuine belief that  $W_a=W_n$  than PB<sub>2</sub>.

---

<sup>25</sup> Stalnaker seems to accept something like this natural understanding of supposition. See his [1987] pp.79-81.

<sup>26</sup> This is a quick statement of my claim. I'll put the point a bit more carefully below, isolating which respects are relevant to successful mimicry. Briefly, those respects are the ones that have to do with our reasoning about what to think, not our reasoning about what to do.

4. When we see what our descriptive reactions are at  $t_1$  while supposing that  $W_a=W_n$ , we're trying to determine what our descriptive reaction would be at  $t_1$  were we to genuinely believe  $W_a=W_n$  at  $t_1$ .

One might here contest one or more of #1-#4. I'll consider that option in some detail in the next section. For now, keep in mind that while there may be different ways of pretending to believe that P, the kind of pretending to believe that is at issue here is one that is to help us understand supposition that a certain scenario is actual. This makes #2 more plausible than it might be otherwise and that makes #3 more plausible by introducing a measure for success in pretend-belief that would be perhaps inappropriate in other contexts.

What does this understanding of supposition that a scenario is actual suggest about the evidential status of our armchair descriptive reactions? To bring out the issue a bit more clearly, compare the Jackson/Chalmers strategy for identifying our descriptive dispositions with a similar strategy for identifying a different, non-descriptive behavioral disposition. Suppose Jack is not certain how much physical courage he possesses. In order to determine to what extent he is disposed to face danger to promote a good end, he and his friend Bob engage in a series of thought-experiments. They are careful to focus their attention not on identifying what Jack thinks he ought to do, in each of the scenarios they consider, but on identifying what he would do. If they could identify what he would do in each of the scenarios they consider, they could identify the pattern in his reactions and so the extent to which he is disposed to face danger.

Fortunately for them both, Bob is a really great story-teller and Jack has acute suppositional powers. Bob describes each of the scenarios Jack is to suppose actual in

great detail and Jack works his best to suppose that things really are as his friend so vividly describes. Could Jack and Bob fail to identify the content of Jack's disposition for a new reason, not because of Jack's inattention or irrationality or Bob's unwillingness to describe a wide variety of cases, but simply because Jack is wrong about how he would react, were things to turn out in one of the ways Bob describes?

It certainly seems as if he could, even if he represents the situation to himself carefully and in detail. No matter how carefully and to what degree of detail we suppose a scenario to be actual, our reaction to that supposition may come apart from our reaction to the belief that we are in that very situation. The moral of this story is that conducting our best tests for discovering our behavioral dispositions requires the triggering of those dispositions. In the case of our disposition to act courageously, that requires that the subject believe and not merely suppose himself to be in a dangerous situation.

To be sure, there are important differences between our disposition to act courageously (a non-descriptive disposition) and our descriptive dispositions. For one thing, we may have no policy, implicit or explicit, guiding our disposition to act courageously. In the case of our disposition to describe things using some expression "E", in contrast, we don't merely find ourselves inclined to use "E"; we seem to be trying to mark a division among ways things can be. Guiding our careful and informed uses of an expression seems to be something like a referential policy.

But the basic issues raised in the case of courage are in important respects parallel to the issues raised in the armchair evaluation of intuition-pumpers. In both cases, we're trying to discover the content of a behavioral disposition D by finding a pattern among D's triggerings. In the courage case, Jack's suppositions that various dangerous

situations are actual are designed to create conditions that mimic the belief that he is in danger. This provides an explanation for why Jack's armchair musings provide highly defeasible data for hypotheses about his disposition to face danger: Jack's armchair reactions are defeasible precisely because armchair suppositions typically don't sufficiently mimic belief. Sufficient for what? Sufficient for triggering Jack's disposition to face danger. For this reason, finding the pattern in Jack's armchair reactions isn't the best way to discover what his disposition is. That's because those reactions aren't typically triggerings of Jack's disposition to act courageously.

Is this also true when the disposition in question is a disposition to use an expression? Does triggering one's disposition to describe things in E-terms require beliefs about how things are? Answering in the affirmative here seems too strong. But why does it seem too strong? It's too strong because of a further difference between the descriptive case and the case of supposing that one is in danger. In the danger case, it's just hard to conjure up the fear that normally accompanies the belief that one is in danger. But, except perhaps in cases in which "E" is an expression for a property or kind that carries strong affective associations,<sup>27</sup> it's plausible to think that supposing or pretending

---

<sup>27</sup> Perhaps "...is disgusting" is such an example. Does this mean that such terms pose a problem for the Jackson/Chalmers project? Here is a place where the differences between their two views matters. Chalmers' arguments for dualism commit him to the claim that we shouldn't be able to work out what, if anything, a fully or partially phenomenal term applies to in a scenario without explicit, phenomenal information, information that will only be available from a first-person perspective. Given that, by his lights, armchair reasoners should need not only descriptions of scenarios, but also phenomenal information about what it's like to be in or experience such scenarios, at least when testing their dispositions to apply such terms. (Thanks to Ryan Wasserman for helpful discussion here.) Is this a problem for Chalmers' argument for a priori knowledge of reference-fixers, at least as applied to our fully or partially phenomenal terms?

The answer here turns on what one can require of ideal armchair deliberation compatible with counting as of yielding a priori knowledge. On one view, such ideal deliberation may idealize beyond the scope of human cognitive limitations. That would allow for the requirement that ideal armchair deliberators possess phenomenal information not actually available to humans a priori. In the context of the Jackson/Chalmers project, this view doesn't seem attractive. The result would be implicit a priori knowledge of extension-fixers that speakers are in principle incapable of manifesting. If our dispositions to apply terms to scenarios

to believe that  $W_a=W_n$  often sufficiently mimics the belief. Sufficient for what?

Sufficient for triggering or near-enough triggering our real descriptive disposition.

Given this, it's easy to see why our armchair reactions can provide some prima facie evidence for referential hypotheses. What we're doing when we evaluate intuition-pumpers from the armchair is pretending to believe that  $W_a=W_n$  with a degree of success sufficient to trigger our disposition to describe things in just the way we would were we to come to believe that  $W_a=W_n$ . Except in atypical cases in which there is a barrier to getting supposition to sufficiently mimic belief, we have prima facie reason to think that our armchair reactions are triggerings or closely resemble triggerings of our real descriptive disposition.

This is why our armchair descriptive reactions can often provide reliable prima facie data for referential hypotheses. Here, as in the courage case, our armchair reactions

---

supposed to be actual is to be evidence for the claim that we have such implicit a priori knowledge, considering the scenarios in the appropriate way needs to be within the scope of our cognitive abilities.

So, appropriate armchair deliberation must be within the scope of normal, human cognitive limitations. What does putting together this constraint with the need for phenomenal information to trigger our dispositions to apply phenomenal terms tell us about what is needed for appropriate armchair deliberation in these cases? I confess that here I don't see how to answer this question in a way friendly to Chalmers' project. The problem is that I don't see how to add phenomenal information to what is otherwise a priori deliberation in the Jackson/Chalmers sense, compatible with maintaining the a priori status of that deliberation. For example, adding the information that there is some scenario supposed to be actual in which I apply the term "disgusting" in that scenario wouldn't help. That information doesn't help settle whether I would apply "disgusting" to that very scenario, given my dispositions as they are.

Alternatively, I could just be given phenomenal information about myself, for example, the disgust that I would feel if I were in a certain scenario, together with the non-phenomenal description of that scenario. But that disgust itself isn't something a priori available in the Jackson/Chalmers sense, since it's an empirical feature of the actual world in their sense. So, it looks like adding that information undermines the a priori status of the deliberation. Perhaps there is an additional, Chalmers-friendly way of handling this issue. I myself don't see what it is.

Do fully or partially phenomenal terms pose a problem for Jackson's argument for a-intensions? In contrast to Chalmers, Jackson is a type-A physicalist. This means that, by his own lights, our implicit knowledge of the reference-fixers for such terms confers upon us a capacity to correctly apply those terms to non-phenomenally characterized scenarios supposed to be actual. Given that our dispositions to apply those terms from the armchair is to serve as evidence for the hypothesis that we have implicit a priori knowledge of a-intensions, the appropriate armchair deliberation must be within the scope of human cognitive limitations. This is just what seems not to be the case. It just doesn't seem possible, from a description of a scenario supposed to be actual, for one to trigger one's genuine disposition to characterize those scenarios as disgusting, or fun, or....

provide such data only insofar as the reactions themselves are triggered by a supposition that  $W_a=W_n$  that closely resembles belief that  $W_a=W_n$ . In both cases, the ultimate court of appeal for such hypotheses is an individual's reactions under conditions of belief.

This makes the evidential gold standard for such hypotheses empirical for two reasons. First, in the Jackson/Chalmers sense and the sense adopted here, a priori warrant for a claim is warrant one is able to have while suspending all belief as to which world is actual. Evidence we can only obtain by believing the world to be a certain way is, by this definition, not a priori, but a posteriori. Since gold-standard evidence for hypotheses about our behavioral dispositions requires belief that  $W_a=W_n$ , gold-standard evidence is a posteriori.

Second, what our reaction will be to a scenario believed to be actual is an empirical feature of the actual world in the sense reserved here. To determine whether a feature  $F$  of the actual world is empirical in the Jackson/Chalmers sense or not, we see whether we can find another world considered as actual that lacks  $F$ . We can see that our reaction to the belief that  $W_n$  is actual is an empirical feature of the actual world in this sense by noting that it is possible to suppose to be actual different worlds over which our reaction to the belief that we are in  $W_n$  varies. This is true regardless of whether a descriptive or a non-descriptive behavioral disposition is in question. I can imagine it's turning out that I will react to believing myself to be in dangerous situation  $D$  by acting bravely and equally well imagine acting cowardly. So, the way that I will react to coming to believe myself to be in  $D$  is an empirical feature of the actual world. The same applies to our descriptive reactions. I can imagine it's turning out that I react to the belief that I am in

a certain world  $W_n$  by describing it in a number of different ways. The way that I will describe it on coming to believe that  $W_a=W_n$  is an empirical feature of the actual world.<sup>28</sup>

To see the contrast between the Jackson/Chalmers explanation of our evaluation of intuition-pumpers and the one advocated here, consider an example. According to Jessica Wilson<sup>29</sup>, there was a point in the history of science when chemists held

(A) All acids contain oxygen

in such a way that, from the armchair, they refused to recognize the possibility of something's being an acid, but failing to contain oxygen. How should we understand what happened when chemists discovered HCl, a substance now recognized as an acid that fails to contain oxygen?

We can imagine a number of different things that might have occurred. Chemists might have realized that their old definition of 'acid' was not as theoretically useful as a new definition that permitted what would henceforth be called 'acids' to be oxygenless. If that is what happened, then the discovery issued in a change in meaning. Alternatively, chemists might have decided that there wasn't really anything that they were referring to with 'acid', having assumed that there were a bunch of properties that always got co-instantiated, when it turns out that there wasn't.

But another, very natural way that events might have unfolded is as follows:

Chemists accepted (A). If asked before the discovery to carefully consider whether something could count as an acid that failed to contain oxygen--in particular, whether

---

<sup>28</sup> A slightly more technical way of putting the point is: Let R be S's actual "E"-descriptive reaction to the belief that she's in  $W_n$ . Whatever R is, we can imagine a world in which her "E"-descriptive reaction is not R. This makes R an empirical feature of the actual world. But then anyone's knowledge that the actual world possesses R is a posteriori in the Jackson/Chalmers sense. This is true for S herself. So S's knowledge of what her "E"-descriptive reaction R is is a posteriori.

<sup>29</sup> Wilson [2006].

something like HCl could--they would have said “no”. But upon learning of the existence of HCl, they found themselves inclined to call it “an acid”, despite its lack of oxygen. Indeed, they may have found themselves newly inclined to say that they recognized HCl as just the kind of thing they meant to be talking about with “acid” all along. The discovery, it seemed to them, showed them that they had been wrong in their earlier, armchair thinking that their referential policy excluded oxygenless acids.

It doesn't really matter which of these cases is the actual one. What's important here is the naturalness of the last case. In this last case, speakers' carefully considered, descriptive reactions to a scenario believed to be actual are different from their armchair reactions. Their armchair reaction is to refuse to recognize anything that failed to contain oxygen as within the extension of their term “acid”. But on coming to believe that the world contained HCl, they found themselves inclined to call it “an acid,” despite its lack of oxygen. Moreover, they found themselves with a certain second-order reaction to both their previous armchair and later, actual reaction. That second-order reaction was to regard their armchair reaction as mistaken and their actual reaction as better reflecting what they had intended to be referring to all along.

We can even provide a plausible background explanation for how just such a case might occur. Suppose that chemists prior to the discovery of HCl were very firmly convinced that there are no oxygenless acids. Borrowing a Quinean metaphor, we might say that this belief was central to their overall web of belief. In such a case, it may be quite difficult—for all intents and purposes, psychologically impossible--to entertain a supposition that *P is actual* when one firmly believes not-P. (A clear case of this kind might be trying to suppose that the actual world is one in which one doesn't exist.) Given

this psychological barrier to the scope of our suppositional powers, it may be that suppositional-attempts in such cases don't—and psychologically can't--sufficiently mimic belief. Sufficient for what? Sufficient to trigger our descriptive dispositions. Such cases are like the non-descriptive case of Jack trying to discover from the armchair his disposition to face danger. (Earlier I said that “except in atypical cases in which there is a barrier to getting supposition to sufficiently mimic belief, we have prima facie reason to think that our armchair reactions are triggerings or closely resemble triggerings of our real descriptive disposition.” This seems to be just such an atypical case.)

In the next section, I'll consider two objections (#1 and #2) to the above characterization of the acid example. But for now, what does the naturalness of such cases suggest? Minimally, it suggests that at least sometimes our full knowledge of our descriptive dispositions and so of our referential policies rests on empirical evidence. But it also suggests something about the evidential status of our armchair descriptive reactions in general. Above we saw some reason often to count such reactions as prima facie, a priori evidence for hypotheses about our descriptive dispositions. But whether such prima facie evidence ultimately misleads us, as it did the chemists in the imagined case, requires knowing what we can't in principle know from the armchair, i.e. what our considered, careful, and informed descriptive and second-order reactions would be on coming to believe that the world has certain empirical features.

Call “our actual, first-order reactions” our descriptive reactions on coming to believe that the world has certain features F. Call “our actual, second-order reactions” our reactions towards both our armchair and actual, first-order reactions on coming to believe that the world is F. Consideration of the imagined “acid” case certainly doesn't

suggest that our actual first- and second-order reactions will typically undermine the evidential status of our armchair, first-order reactions. If it did, the descriptive case would be more like the courage case and there would be little reason to regard our armchair descriptive reactions as even prima facie evidence for hypotheses about the contents of our descriptive dispositions and referential policies. Rather, it suggests that we can't know from the armchair whether our armchair descriptive reactions are any more than good, prima facie evidence for such hypotheses. That's because:

1. Our ideally rational, armchair supposition that  $W_a=W_n$  aims to mimic the belief that  $W_a=W_n$ , at least in the respects relevant to the issue at hand.
2. Knowing whether our ideal suppositions succeed in this regard requires knowing
  - a) those considered, careful, armchair reactions to those suppositions are the same as our considered, careful, actual, first-order reactions or b) at least that our actual, second-order reactions are not to regard our actual, first-order reactions as undermining or supplanting the evidential value of our armchair reactions.
3. Knowing either 2a or 2b requires knowing whether or not the actual world has certain empirical features, namely, the actual, undermining reactions in question.
4. Given #3, knowing whether our armchair descriptive reactions carry more than prima facie evidential weight requires knowing whether or not the actual world has certain empirical features.
5. Therefore, knowing whether our armchair reactions carry more than prima facie evidential weight is a posteriori knowledge.<sup>30</sup>

---

<sup>30</sup> One might think here that this discussion is grist for the Jackson/Chalmers mill. After all, doesn't it show that I am able to know a priori what it takes for something to be within the extension of my term, e.g. "water"? I seem to know this after all: I know that in any world  $W_n$  supposed to be actual, what I am disposed to call "water" in  $W_n$  fixes the extension of "water" in  $W_n$ . (Thanks to David Braddon-Mitchell

What are the respects in which the suppositions at issue in intuition-pumpers must mimic beliefs? Here is one respect that is clearly inappropriate: One role that beliefs play in our reasoning is to help us figure out what it makes sense to do. If, on a hot day, I want an Italian ice, my beliefs about the location of the nearest vendor of good ice become relevant to my deliberations. But I can not only suppose things I'm agnostic about, often I can suppose things I positively disbelieve. For example, I positively disbelieve that my local sporting goods store is the nearest vendor of good ice. But I can suppose that it is. If I allowed my supposition to play the same role in my reasoning about what to do that a belief with the same content would play were I to have that belief, I may well find myself looking for ice on a hot day at my local sporting goods store.

But this difference between suppositions and beliefs is irrelevant in the present context. The role that suppositions play in our evaluations of intuition-pumpers is in getting us to reason about what to think, in particular, what to think about what counts as what, supposing that the actual world is a certain way. So we need merely limit the scope of the respects in which our suppositions are to mimic our beliefs to the respects that have to do with our reasoning about what to think in order to get claim #1 to come out true.

Do considerations #1-5 show that we *can't* have a priori knowledge of the extension-fixers for our terms? Here the answer depends in part upon our resolution of the two issues

---

for pressing this point in conversation.) But, 1) this reply isn't available to either Jackson or Chalmers neither of whom wants the extension-fixers to be features of the worlds considered as actual since this would make metalinguistic facts about those worlds relevant to their evaluation. (Chalmers [2006]). Moreover, 2) making room for interesting a priori truths requires that such metalinguistic facts are irrelevant to their evaluation. If they were relevant, every expression would have a trivial and different a priori known component of content, since it would be true for every expression "E" that, for every world  $W_n$ , what "E" refers in  $W_n$  to has the feature of being what I call "E" in  $W_n$ . But then, for any two terms "E" and "T" we take, there will be some worlds we may consider—either suppose or believe—to be actual such that I will call different things "E" and "T" in those worlds. *That* is why Jackson and Chalmers must insist that metalinguistic facts that obtain in a world of evaluation are irrelevant to their evaluation. Allowing their relevance requires giving up on interesting, a priori truths.

discussed in the introduction. That first issue is about what justificatory standard is required for knowledge. If knowing that P requires gold-standard evidence that P, then the answer is “yes”. If knowledge requires no more than good, prima facie evidence, then the answer *may* be “no”. Just as in the a posteriori case, where some justification, plus belief, plus truth isn’t enough for knowledge, so too in the a priori case. It depends at least upon how strong a believer’s justification is.<sup>31</sup>

The second issue has to do with whether to hold that properly a priori knowledge could have both an empirical defeater and an empirical falsifier. Earlier I said that this issue is at least partly terminological. There are a couple of reasons to resolve this issue by treating our knowledge of our referential policies as a posteriori. For brevity’s sake, call such knowledge “referential knowledge”. These reasons have to do with differences between referential knowledge and classic cases of putative a priori knowledge. I’m here agnostic about whether there is any knowledge that is a priori in any attractive sense. But I accept two plausible constraints on any attractive account. The first is that those classic cases should come out a priori on any plausible account. Second, grounds for treating contested cases as cases of a priori knowledge should be their similarity to the classic examples.

Two classic kinds of examples of a priori knowledge are definitions or stipulations and mathematical truths. Famous examples of definitions or stipulations include “bachelors are unmarried men” and “Julius is the single inventor, if there is one, of the zip”. Holding fixed the meanings of these expressions (including the stipulation that “Julius” refers to the single individual, if any, who invented the zip), it is hard to imagine either having an empirical defeater. It is easier to imagine empirical defeaters in the case of mathematical claims. After working it out in my head, I may think I know that the sum of a lengthy column of largish

---

<sup>31</sup> Thanks to Albert Casullo for helpful discussion here.

numbers is 11,562, but learn from my calculator that I have made an error. In this case, my knowledge of the correct answer is justified by my empirical knowledge that my calculator gives a different sum, together with my empirical knowledge that my calculator is reliable. In this kind of case, the defeater, a mistake in calculation, is in principle discoverable from the armchair. We can also imagine cases in which this is not so, cases in which what undermines one's a priori justification is only discoverable empirically. For example, we can imagine a case in which mathematical experts falsely tell one that one's armchair calculation contains a mistake. Since there is no mistake, there is no correcting the calculation from the armchair. Nonetheless, arguably, the status of the mathematicians as experts suffices to undermine one's a priori justification. If so, then, since their counter-claim is empirical, knowledge of that claim is itself only available a posteriori.<sup>32</sup> That makes their counter-claim an empirical defeater in principle unavailable from the armchair.

In these respects, mathematical knowledge is similar to referential knowledge as I have characterized it. In both cases, one's a priori justification may be undermined by a possible defeater only empirically available. Does this mean that referential belief prima facie justified by armchair reasoning can be happily assimilated to mathematical belief, prima facie justified by armchair reasoning? If so, should we conclude from this that both are cases of a priori knowledge when true? There is good reason to answer each of these questions in the negative. At this point in my argument, I am considering what follows if my rival explanation of our evaluations of intuition-pumpers is correct. If that explanation is correct, then some a priori justified referential hypotheses are falsified by our actual descriptive reactions. Since our actual descriptive reactions are empirical features of the world in the Jackson/Chalmers sense, these falsifiers are empirical and hence, discoverable only empirically. In contrast, while our a

---

<sup>32</sup> Thanks to Jonathan Schaffer for this example.

priori justification for our mathematical belief may have defeaters that are only knowable a posteriori, they don't have falsifiers that are only so knowable. If mathematical experts falsely tell us that we've made a mistake in our armchair calculation, that doesn't falsify our mathematical belief—only a wrong answer would do that. Wrong answers, though, are in principle a priori discoverable. In this respect, mathematical mistakes are like the mistakes in armchair reasoning that I earlier identified as compatible with the Jackson/Chalmers project. That one has hastily evaluated an intuition-pumper or failed to consider a sufficient variety of such conditionals is something one may discover either empirically or from the armchair. In the latter kind of case, more careful armchair reaction is the falsifier of a referential hypothesis developed on the basis of faulty armchair reactions. These are armchair falsifiers.

If the considerations raised in this section are correct, this contrasts with the possible empirical falsifiers for putative armchair referential knowledge that I am concerned with here. When they are falsifiers of referential hypotheses developed from the armchair, our actual first- and second-order reactions are empirical falsifiers. Moreover, the only way to check for the presence of this kind of falsifier is empirically. There is, then, at least one important in principle difference in kind between the kind of a priori warrant that can be had for our referential beliefs and for our beliefs in the classic cases of putative a priori knowledge. Even in classic cases in which empirical defeaters are always in principle possible (as in the above mathematical cases), empirical falsifiers are not. This is an important difference between the referential case and the classic cases of a priori knowledge. And this difference constitutes some reason not to classify our knowledge of our referential policies together with the standard examples of a priori knowledge.

In sum, there is a range of plausible positions on these two issues. On only a proper subset of those positions does referential knowledge come out a priori. This suffices to show that even if we accept the three Jackson/Chalmers assumptions identified in the introduction, we are not *forced* to conclude that we have implicit a priori knowledge of our extension-fixers in order to vindicate the evidential value of our confidant evaluations of intuition-pumpers.

On the rival explanation proposed here, we can have some prima facie reliable, a priori evidence for referential hypotheses in the form of prima facie warranted a priori hypotheses about our descriptive dispositions. That prima facie warrant is provided by our armchair evaluations of intuition-pumpers, so long as the conditions for thinking them triggerings of our genuine descriptive disposition obtain. Those armchair evaluations are prima facie good evidence precisely because they typically track what our actual reactions would be were we to believe the world to be as supposed.

But prima facie warranted armchair hypotheses have possible empirical falsifiers. Since our best evidence for a claim includes ruling out possible falsifiers, our best evidence for our referential hypotheses is at least partly empirical. This means we needn't think of ourselves as having implicit, a priori knowledge of what property or properties a particular must have in order to fall within the extension of an expression in order to regard our armchair intuitions as providing good evidence for hypotheses about what we're talking about with our use of terms of philosophical interest, such as "good" or "physical". That we must so assume is a central claim in the Jackson/Chalmers argument for a priori, extension-fixing components of content. So, without additional argument, the central Jackson/Chalmers argument for a priori knowledge of an extension-fixing component of content fails.

#### IV. Objections and Replies

In the previous section I noted that it was possible to contest the sketched understanding of supposing a world to be actual. It's time to consider this and other ways of blocking the conclusions of the previous section.

*Objection 1:* More carefully described, the acid case is grist for the a prioricist's mill. In the case as it is described above, the scenario believed to be actual is more detailed and concrete than the scenario supposed to be actual from the armchair. A proper test of whether a speaker's actual and armchair reactions could be distinct without a change in meaning requires that the two scenarios considered have the same content. In the above example, chemists came to believe that the world contained HCl. Properly testing the a prioricist's claim requires that we ask what chemists might have called HCl, had they supposed, but not believed, it to exist.

There are two possibilities here, neither of which supports the conclusion that gold standard evidence for hypotheses about our referential dispositions and policies is a posteriori. The first possibility is that had chemists supposed that the world turned out to contain HCl, they would have described that as it's turning out that the world contained an oxygenless acid. In that case, their first-order actual and armchair reactions would have been the same. So, there wouldn't have been a second-order, actual reaction that their first-order, armchair reaction was mistaken. The actual, first-order reaction, then, would have told them no more than the armchair reaction did about what their referential dispositions and policies were. So, properly run, the armchair experiment would have sufficed for chemists to see that when they held that all acids contained oxygen, they held it as an empirical generalization, not as a condition on acidity.

A second possibility is that the armchair and actual reactions would have been different, for example, that from the armchair, chemists would not have described HCl as “an acid”, but on coming to believe that the world contained it, would have. This case looks like a change in meaning. But the question the a prioricist gives an answer to isn’t: What would we mean were we to be changed in certain ways (e.g. come to believe various things that we don’t now believe)? Rather, the question is: what do we mean with our expressions *now* and how might we know what we mean? In neither of these two possibilities does the actual reaction help answer the latter question.<sup>33</sup>

*Reply 1:* There are several things to say here.

a(i) For this objection to have force, some reason needs to be given for saying why a difference in actual and armchair reactions would indicate a change in meaning. After all, we wouldn’t say a difference in reaction in the courage case meant one’s disposition to face danger had changed. Rather, we’d say that we were wrong about what our disposition is. So, we need some reason to think that were our actual and armchair first-order reactions to differ, the explanation would need to be a change in meaning.

a(ii) If it is reasonable to suppose that coming to believe something we don’t now believe might lead to a shift in what we mean by some expression, we might equally worry that supposing things we don’t now suppose might have the same effect. We might, for example, equally worry that had chemists supposed that the world contained HCl, that would have changed them in a way that changed what they meant by “acid”. For objection 1 to be an objection against the present position, as opposed to a general puzzle for any position in this area, there needs to be a plausible story about what why coming to

---

<sup>33</sup> Thanks to David Braddon-Mitchell for posing this objection in conversation.

believe that P might constitute a meaning-change inducing change in us, while coming to suppose that P does not.

a(iii) No doubt, sometimes supposing things we hadn't previously supposed and believing things we hadn't previously believed causes us to change our referential policies. But that either should always do so would make our descriptive dispositions extremely odd. If they did, then every time we come to believe or suppose something new, our descriptive dispositions, and so our referential policy, would change. But if that is so, then it is very doubtful that those dispositions have the sufficient stability to be properly thought of as guided by a policy. If, as assumed here, our referential policies play an important role in determining the extensions of our expressions, then it would be doubtful that there is anything we talking about with our disjointed, ever-changing use of our expressions.

I take this to be a *reductio* of the worry that our descriptive dispositions fluctuate in the way described.

b) Suppose, though, that there is some reason to think that beliefs are more likely to induce changes in us that lead to changes in what we mean than suppositions do.

Objection 1 rests on an accidental feature of the acid example, namely, that it involves a time-shift. We might ask a speaker how she is now inclined to describe various scenarios, some which she merely supposes to obtain, and others which she believes to. For the reasons given earlier, our descriptive reactions to scenarios believed to be actual are empirical features of the world and our knowledge of what they are is a posteriori. In this case, we don't come to believe what we previously didn't, so there can be no meaning-change inducing change involved—there is no change involved. Here we can

turn the challenge back on the challenger: Why do our armchair, descriptive reactions to our present suppositions have superior evidential value or carry greater evidential weight than our actual reactions to our current beliefs?

c) If we gloss supposition in the way we have in section III, we'll have an argument that our reactions to our suppositions are evidence of our dispositions only insofar as they match what our reactions would be were we to believe at the time of the supposing (hence time indexing in the original formulation of the claim in section III). This is some reason to accord our actual, descriptive reactions to cases believed to be actual greater evidential weight when deciding between different hypotheses as to our descriptive dispositions and policies.

*Objection 2:* The above characterization of the imagined acid example overlooks an important feature of the Jackson/Chalmers position that undermines the lesson drawn. As noted, not any old armchair consideration of an intuition-pumper suffices to trigger one's descriptive disposition. What hasty evaluations and failure to consider a sufficiently large number of intuition-pumpers have in common is that they are both departures from ideally rational armchair reflection. What the acid example shows, at most, is that there is an additional way in which armchair reasoning can fail to reflect that ideal. In the imagined acid example, chemists faced a barrier to the scope of their suppositional powers—they were unable to adequately suppose to be actual a world in which some acids failed to contain oxygen. Had they not failed in this respect, they would have recognized HCl as an acid. Thus, their armchair reaction would not have differed from

(or, more importantly, been evidentially undercut by) their actual reaction. This is all in keeping with the Jackson/Chalmers proposal.<sup>34</sup>

*Reply 2:* Three issues are important here. One is what armchair reasoning must be like in order to count as ideal. A second, related issue is whether the limit on chemists' suppositional powers in the imagined case need be understood as the source of a genuine mistake in armchair reasoning in the sense that it is something that further armchair reasoning could correct. (In other words, whether that limitation must be understood as the falling short of the ideal or whether it may be a limitation on the suppositional powers of even an ideal reasoner.) A third issue is how, exactly, we are to understand the idea of supposing a world to be actual. I'll first lay out how I think these issues should be addressed and then show how that material allows for a reply to this objection.

I'll take the third issue first. The notion of considering a world as actual is central to the Jackson/Chalmers proposal. Unfortunately, neither has very much to say by way of explaining what this central notion amounts to. In section III, I provided a partial elucidation of this notion in terms of a natural understanding of the attitude of supposing that *P*. This elucidation is only partial, however, in leaving dark a parameter in the Jackson/Chalmers notion, namely, that of supposing that *P is actual*.

A full treatment of this topic is well beyond the scope of this paper. Fortunately, replying to the present objection does not require a full treatment. Here I rely on the plausible understanding of that notion favored by Stalnaker.<sup>35</sup> On Stalnaker's way of doing things, considering (what I call 'supposing') a world *w* to be actual involves thinking of *w* as the very world that one is in. This understanding has two advantages.

---

<sup>34</sup> Thanks to Jill North, Joshua Schechter, and Jessica Wilson for each raising something like this objection.

<sup>35</sup> Stalnaker [2003] p.196.

First, if we accept Stalnaker's way of thinking of this notion, we'll have a notion that fits with the standard way of marking that some world is being supposed to be actual as opposed to supposed to be counterfactual. Standardly, we mark this distinction with the locutions "if it turns out that P..." and "if it had been that P...". In the case of the former, we're considering the P-world as actual, in the latter, as counterfactual. "If it turns out that P" suggests that we think of P as the world we turn out to be in. A second advantage of this understanding is its fit with the standard, intuitive understanding of actuality. On that intuitive notion, the actual world is this very one that one finds oneself in.<sup>36</sup> "Actually" seems to function as a worldly "I am here" operator. I've said that this understanding is plausible, but let me now go on record with a claim that is a bit stronger: If we reject Stalnaker's notion, I don't see what notion we could put in its place that avoids the abuse of the English language.

So much for the third issue. What about the first? Neither Jackson nor Chalmers says very much about this issue. Let me here take a stand, one that I believe all parties to this dispute should accept. In a case such as the present one where part of the issue is what, if anything, to count as implicit a priori knowledge, the armchair reasoning by which some truth is discoverable should be within the cognitive powers of the agent to whom we wish to attribute such knowledge. Remember, the Jackson and Chalmers position is that our dispositions to apply our terms to scenarios supposed to be actual are manifestations of our implicit, a priori, referential knowledge—that's why our evaluations of intuition-pumpers are evidence for the hypothesis that we have such knowledge. To be able to manifest such knowledge by evaluating intuition-pumpers,

---

<sup>36</sup> Jackson himself in places favors this understanding. See his [1998a] p. 48. For another fan of a priori extension-fixers who also seems to favor this understanding, see Braddon-Mitchell [2004] p. 137.

though, we must be able to suppose their antecedents to be actual in the right way. That requires that such suppositions are within the scope of our cognitive abilities. The ideal armchair reasoning that undergirds claims to implicit a priori knowledge must be cut down to human size. Any plausible resolution of the first issue, I submit, must accommodate this.

We arrive, finally, at the second issue. Let me just baldly state what I take to be a very plausible hypothesis about human psychology: In cases in which we've got a very, very firmly entrenched empirical belief, it is sometimes psychologically impossible for us to suppose that the actual world is one in which that belief is false. Embedded in my discussion of the acid example, I offered what may be the paradigm case of this kind: It is an empirical belief of mine that I exist.<sup>37</sup> But I just can't suppose that it turns out that the actual world is one in which I don't exist. Putting it a slightly different way, I just can't get myself to suppose that it may turn out that I don't exist.

So far we have (i) the ideal armchair reasoning that is required for a priori knowledge must be ideal given-the-psychology-of-the-knower-in-question and (ii) where humans are the knowers at issue, (i) means that ideal reasoning cannot require that we perform psychological impossibilities, (iii) at least some of our empirical beliefs are such that supposing that they are false of the actual world is psychologically impossible. This gets me part of the way towards the conclusion I need, but not all of the way. To get all of the way, I need one more plausible claim about human psychology.

The belief that I exist is special in this sense: It is a belief, as Descartes points out, that I cannot fail to have. Given this, it differs in a crucial respect from the imagined acid

---

<sup>37</sup> I suppose that one might argue that "I exist" is a priori known by me. However, Chalmers regards such truths as empirical. See his [2004].

case. In the imagined acid case, the deeply entrenched empirical belief is one that I psychologically could come to reject. The imagined acid case, then, is a case in which the entrenchment of an empirical belief is so deep that a thinker is unable to seriously suppose that it turns out to be false, but which she nonetheless comes to reject.

This may at first seem like a surprising case. But on second thought it shouldn't be. What we have here is someone who has, in Stalnaker's helpful terms, narrowed down the set of possibilities to a set that excludes the truth of the supposition. When the possibilities ruled out are empirical, such a thinker's attitude is that certain things *just aren't going to happen*. We can imagine how that may be so for a chemist in the imagined scenario. Given her entrenched empirical belief in certain chemical theories, she regards the possibility that the actual world contains a chemical compound that fails to contain oxygen yet behaves just like an acid as just not a live option. It needn't be that she can't conceive of such a compound. Rather, what she can't conceive of is *the actual world* being one that contains such a thing. For her, all such worlds are counterfactuals, worlds in which (as she would put it) the laws of chemistry differ from the actual laws. But this is a state of mind that can be changed.

Now, helping ourselves to Stalnaker's understanding of the notion of supposing a world to be actual, we can put the point more broadly. Take a case where I am certain that P. So, there are no possibilities in the set of possibilities that make up my belief state in which  $\sim P$ .<sup>38</sup> On the present understanding, supposing that  $\sim P$  is actual requires that I ask myself to suppose that it turns out that  $\sim P$ . Which possibilities am I to entertain when I engage in this supposition? None that is compatible with my current belief state. Given

---

<sup>38</sup> This will be true on Stalnaker's way of representing belief states which I am adopting for present purposes. See his [1987] p. 69.

the much-discussed difficulties with believing and, we might add, disbelieving at will, we've arrived at a type of supposition that I psychologically cannot take up willy-nilly. In the current kind of case, supposing that  $\sim P$  is actual requires coming to believe it.<sup>39</sup>

*Objection 3:* Supposing that  $W_a=W_n$  is not "a kind of pretend belief that  $W_a=W_n$ " as the argument in section III assumes.

*Reply 3:*

a) We may think of my initial gloss on supposition as a kind of pretend-belief as an intermediary step along the way to the claim that in the sense at issue in the evaluation of intuition-pumpers, successful suppositions mimic beliefs in the relevant respects. As long as that claim is conceded, the intermediary step isn't crucial. And that we do allow suppositions to figure in our reasoning about what to think in the way that beliefs do constitutes some good reason to think that successful suppositions mimic the role of beliefs in our reasoning in these respects.

b) In order to evaluate the Jackson/Chalmers argument summarized in sections I and II, we need to know what is involved in supposing a certain world to be actual. So, if we reject understanding supposing a world to be actual as pretending to believe that it is, we'll need some other understanding to replace it. The challenge for the fan of a priori-extension-fixers will then be to say why, given the favored understanding of the suppositions in question, our reactions to those suppositions have a kind of superior evidential value for referential hypotheses than do our reactions to beliefs with the same content.

---

<sup>39</sup> Edgington makes a somewhat similar point in her [2006] p.9-10. There she notes that the suppositional theory of indicative conditionals requires that the probability of the antecedent not be zero, that is, an evaluator must regard the antecedent as an epistemic possibility.

*Objection 4:* The suppositions in question can be understood as pretend-beliefs, but “S’s pretending to believe that  $W_a=W_n$  at time  $t_1$ ” does not aim “to mimic the state that S would be in were S to believe that  $W_a=W_n$  at  $t_1$ ”.

*Reply 4:* Here, the challenge for the fan of a priori-extension-fixers will be to say what kind of pretend-belief illuminates supposition, but does not aim to mimic belief. In addition, a deflationary explanation will need to be given for what looks like evidence in favor of the natural understanding of supposition. If suppositions aim to mimic belief, that would explain why we normally allow them to figure in our reasoning in the way genuine beliefs do, at least as regards the role of belief in reasoning about what to think and for the duration of our supposing. Why would we do this, if suppositions don’t aim to mimic genuine beliefs in these respects?

And, again, that fan will need to say why, given her favored understanding of pretend-belief, our reactions to those pretend-beliefs have a kind of superior evidential value for hypotheses about our referential policies than do our reactions to beliefs with the same content.

## **V The Role of Intuitions about Possible Cases**

On the Jackson/Chalmers proposal, our armchair evaluations of intuition-pumpers allow us to make explicit our implicit a priori knowledge of an extension-fixing component of content. Our grasp of these extension-fixers is the grasp of what property or set of properties something must have to fall within the extension of the associated terms.

Also, on that proposal, such extension-fixers may be thought of as concepts associated with expressions. Conceptual analysis in their sense just is the process of

making explicit our implicit, a priori knowledge of our extension-fixers by the armchair evaluation intuition-pumpers. That, according to Jackson, is just what we're doing when we consider philosophical thought experiments like Gettier's and Putnam's.<sup>40</sup>

According to Jackson and Chalmers, such analyses play a necessary role in the resolution of metaphysical disputes in the philosophy of mind, epistemology, and metaethics by providing a kind of job description for the targets of those disputes. Their idea is that using our armchair evaluations of intuition-pumping conditionals, we can in principle identify what property or set of properties something has to have in order to count as an instance of the kind that is the target for reduction. Once we've identified what those properties are, we need only look in our candidate reduction base for what has those properties and we'll have found the reduction base for our target. This is their reasoning behind their much-discussed claim that conceptual analysis is required to resolve metaphysical disputes in a variety of areas, including the dispute between physicalist and dualists about the nature of the phenomenal.<sup>41</sup>

Do the considerations raised above show that our armchair judgments about the truth-value of intuition-pumping conditionals have no role to play in the resolution of such disputes? Do they show that conceptual analyses in the Jackson/Chalmers sense have no role to play? I think the answer in both cases is "no". They do, however, show that conceptual analysis is not the only form our explicit knowledge of our extension-fixers may take. Our knowledge of our extension-fixers may be a posteriori. Conceptual analyses, in contrast, are, by definition, a priori. Given this, these considerations do show how we might have a posteriori knowledge of what property or set of properties

---

<sup>40</sup> See Jackson [1998a].

<sup>41</sup> See, for example, Chalmers and Jackson [2001] and Jackson [1998a].

something has to have in order to count as an instance of the kind that is the target for reduction. That's the role that Jackson and Chalmers argue can only be played by analyses. So, these considerations do show how we might have reductions vindicated by their method, but without analyses.<sup>42</sup>

Consider the example of the dispute over the content of our concept of the basically physical, as it figures in the thesis of physicalism. The example is a useful one, since it provides illustrations both of the kind of a priori evidence for the content of our concept that Jackson and Chalmers regard as crucial and the empirical kind that I am arguing is superior. When physicalists claim that there is nothing over and above the physical, they do not intend for 'the physical' to be understood in an entirely new, stipulated sense, but to bear significant continuity with our pretheoretical notion. One way of understanding what is at issue between the different accounts of the physical<sup>43</sup> is over how best to understand what is central in that pretheoretical notion for the purposes of formulating the thesis of physicalism. On the method for discovering the content of our pretheoretical notions I am here recommending, consideration of our intuitions about hypothetical ways the world might turn out to be provides empirically defeasible, a priori evidence for and against accounts of the content of the pretheoretical concepts in play in those intuitions. To illustrate: Suppose we found we have a widely-shared intuition that if our best-confirmed physical theory were to turn out to posit something that lacked mass, that would be its turning out that being physical does not require having mass. On the view I'm defending, that armchair intuition would provide empirically defeasible, a priori evidence that our concept of the physical has two features. The first is that our concept is

---

<sup>42</sup> For a more detailed defense of the claim that analyses are not necessary for reductions to be vindicated by what is otherwise the Jackson/Chalmers method, see Dowell [forthcoming].

<sup>43</sup> See, for example, Dowell [2006] and Wilson [2006].

such that being physical does not require having mass. The second is that the content of our concept is tied to what our best-confirmed physical theory tells us there is.

When it does turn out that our best-confirmed physics does posit forces and particles lacking mass and we continue to count such forces and particles as physical, we're given some empirical confirmation of both of our armchair hypotheses about the content of our concept of the physical, namely, that indeed, being physical does not require having mass and the content of our concept is tied to what our best physical theory tells us there is. This makes the ultimate justification of our armchair hypothesis about the contents of our concept of the physical empirical and so our knowledge of its content a posteriori. The acid example is a reverse case, a case in which our armchair intuition is falsified by our actual reaction.

Let's grant, as seems plausible, that Jackson and Chalmers are right to hold that knowledge of what it takes to be a T, according to our pretheoretical notion of Ts, has an important role to play in settling disputes about the nature of Ts. The considerations raised here suggest that the role that intuitions about hypothetical cases play in the Jackson and Chalmers account is better played by our intuitions about how to describe cases we believe are actual. Sometimes, however, the world doesn't provide the cases that we may be interested in discovering our actual reactions to, for example, how we would be prepared to describe a highly idealized scenario. In such cases, discovering our armchair reactions to scenarios supposed to be actual is the closest to discovering our actual reactions we're going to get. Because it is reasonable to suppose that (unlike in the courage case) our armchair reactions at least often match our actual reactions, it is reasonable to treat those armchair reactions as good, prima facie evidence for referential

hypotheses. In cases in which we can't get ourselves to believe in the actuality of scenarios we'd like to test, our armchair hypotheses provide the best evidence we'll have.

So, the considerations from section III suggest that while conceptual analysis isn't always necessary for the resolution of metaphysical disputes,<sup>44</sup> as Jackson and Chalmers have argued<sup>45</sup>, it is very often useful.

## **VI. Conclusion**

There has been a lively discussion in the recent literature in the philosophy of language on the two dimensional framework over whether there is good reason to think that competent speakers have a priori knowledge of the extension-fixers for many of their terms.<sup>46</sup> The greatest challenge facing the foe of a priori extension-fixers is to provide an explanation of our confident evaluations of intuition-pumpers that vindicates their evidential role in helping to settle philosophical disputes. I know of no attempt among such foes to meet this challenge. I hope to have offered one way of meeting it here.

There has also been a lively discussion of the role such extension-fixers might play in the literature on resolving metaphysical disputes, particularly disputes about the metaphysics of mind.<sup>47</sup> Central to Jackson's arguments for type-A physicalism<sup>48</sup> and Chalmers' arguments for dualism<sup>49</sup> is their claim that conceptual analysis is required for the vindications of metaphysical reductions, e.g. of pain to some physical state. Since for them, conceptual analyses are just what we get when our implicit a priori knowledge of our extension-fixers is made explicit, this means that, for them, the very possibility of

---

<sup>44</sup> For a more detailed defense of the claim that conceptual analysis isn't needed to defend the reductions at issue in many metaphysical disputes, see Dowell [forthcoming].

<sup>45</sup> See Jackson and Chalmers [2001].

<sup>46</sup> See, for example, Byrne and Pryor [2006], Soames [2005], and Yablo [2006].

<sup>47</sup> See, for example, Jackson and Chalmers [2001], Jackson [1998], Block and Stalnaker [1999], McLaughlin [2005].

<sup>48</sup> Jackson [1998a].

<sup>49</sup> Chalmers [2002].

vindicating such reductions requires a priori extension-fixers. Elsewhere I have argued that conceptual analyses are not required for vindicating reductions by giving examples of such vindications that don't require a priori knowledge of extension-fixers. There I argued that a posteriori knowledge of our extension-fixers must be in principle possible, but did not provide a method for its acquisition.<sup>50</sup> Here I hope to have strengthened that case by identifying a procedure for the empirical vindication of referential hypotheses.

## References

- Block, Ned and Robert Stalnaker 1999. "Conceptual Analysis and the Explanatory Gap", *Philosophical Review* 108: 1-46.
- Braddon-Mitchell 2004 "Master's of Our Meanings", *Philosophical Studies* 118: 133-152.
- Byrne, Alex and Jim Pryor 2006. "Bad Intensions", Garcia-Carpintero, Manuel and Josep Macia (Editors) *Two-Dimensional Semantics*, Oxford: Oxford.
- Chalmers, David 2002a. "Consciousness and It's Place in Nature", .  
*Philosophy of Mind: Classical and Contemporary Readings*, D. Chalmers, ed., New York: Oxford: 247-72.
- Chalmers, David 2002b. On Sense and Intension, *Philosophical Perspectives* 16: Mind and Language.
- Chalmers, David 2004. "Two-dimensional Semantics", *Philosophical Studies* 118: 153-226.
- Chalmers, David 2006 "The Foundations of Two-Dimensional Semantics", Garcia-Carpintero, Manuel and Josep Macia (Editors) *Two-Dimensional Semantics*, Oxford.

---

<sup>50</sup> Dowell [forthcoming].

- Dowell, J.L. 2006. "The Physical: Empirical, Not Metaphysical", *Philosophical Studies* 131: 25-60.
- Dowell, J.L. *forthcoming*. "Serious Metaphysics and the Vindication of Reductions", *Philosophical Studies*.
- Edgington, Dorothy 2006. "Conditionals", <http://plato.stanford.edu/entries/conditionals/>
- Frankfurt, Harry 1969. "Alternate Possibilities and Moral Responsibility", *Journal of Philosophy* 66: 829-839.
- Gettier, Edmund 1963. "Is Justified True Belief Knowledge?", *Analysis* 23: 121-3.
- Garcia-Carpintero, Manuel and Josep Macia 2006. (Editors) *Two-Dimensional Semantics*, Oxford.
- Jackson, Frank 1998a. *From Metaphysics to Ethic: A Defence of Conceptual Analysis*, Oxford: Clarendon.
- Jackson, Frank 1998b. "Reference and Description Revisited", *Philosophical Perspectives* 12: 201-18.
- Jackson, Frank 2001. "Locke-ing Onto Content", *Royal Institute of Philosophy Supplement* 49: 127-43.
- Jackson, Frank and David Chalmers 2001. "Conceptual Analysis and Reductive Explanation", *Philosophical Review* 110: 315-61.
- Kripke, Saul 1980. *Naming and Necessity*, Cambridge: Harvard.
- McLaughlin, Brian. 2005. "A Priori versus A Posteriori Physicalism." In Christian Nimtz and Ansgar Beckermann, eds., *Philosophy-Science -Scientific Philosophy*, Main Lectures and Colloquia of GAP 5, Fifth International Congress of the Society for Analytical Philosophy, 2003. Paderborn: Mentis, pp.267-285.

- Salmon, Wesley 1993. "Causality: Production and Propagation" in Ernest Sosa and Michael Tooley (Editors) *Causation*, New York: Oxford University Press: 154-171.
- Schroeter, Laura and John Bigelow *forthcoming*. "Jackson's Classical Model of Meaning," Ian Ravenscroft (editor), *Minds, Worlds, and Conditionals: Themes from Frank Jackson*, Oxford: Oxford University Press.
- Soames, Scott 2005 *Reference and Description: The Case Against Two-Dimensionalism*, Princeton: Princeton.
- Stalnaker, Robert 1987. *Inquiry*. Cambridge: MIT Press.
- Stalnaker, Robert 2003. "On Considering a Possible World as Actual", *Ways a World Might Be*, Oxford: Clarendon:, pp.188-200.
- Thomson, Judith Jarvis 1985. "The Trolley Problem", *The Yale Law Journal* 94.
- Weatherson, Brian 2003. "What Good Are Counterexamples?", *Philosophical Studies* 115: 1-31.
- Wilson, Jessica 2006. "On Characterizing the Physical", *Philosophical Studies* 131 vol. 1: 61-99.
- Yablo, Stephen 2000. "Textbook Kripkeanism and the Open Texture of Concepts", *Pacific Philosophical Quarterly* 81: 98-122.
- Yablo, Stephen 2005. "No Fool's Cold", Garcia-Carpintero, Manuel and Josep Macia (Editors) *Two-Dimensional Semantics* , Oxford: Oxford.