

n -Dimensionalism

Philipp Keller*

June 29, 2001

Draft
comments welcome

Abstract

Physicalistically inclined readers of Chalmers are likely to wonder: Even if Zombies are conceivable, this fact (if it is one) does not seem to show much about *our world*, which does not contain any (or so we suppose)? Physicalism may still be contingently true.¹ Here Two-Dimensionalism enters the stage, specifying conditions under which the step from conceivability to possibility is justified. I will argue, however, that Two-Dimensionalism does not help to dissipate this first of all objections, for it turns crucially on a distinction which can and has to be applied across the board. If we are willing to make it, we get n -Dimensionalism, for any n . Stopping with $n = 2$ begs the question against the physicalist.

Whether or not you can conceive that p depends on your epistemic situation. Your epistemic situation depends on the world you are in. This much is uncontroversial. What is more important is that your epistemic situation also includes what you *know* or even what you *are in a position to know* about the world you are in. Not all of this knowledge can be stripped off while determining truth or falsity of this-worldly statements in counterfactual circumstances. I want to claim that this trivial fact undermines the distinction between *considering a possible world w as counterfactual* and *considering w as actual* and the associated distinction between primary and secondary intensions. The problem turns on the question what it means to 'keep our language (or our concepts) fixed' when considering other worlds. To "retain the concept of the *real* actual world" (CC: 8, 42), we have to make assumptions about how the real world is like. These assumptions, I will argue, trivialize what we can find out about what is a priori in this way.

1 The illusion of contingency

"Water is H_2O ", Kripke told us, is necessary. At first, this seems counterintuitive. Are we not able to imagine possible worlds where it is false; is not Putnam's Twin Earth a world in case? It is not, says Kripke, for by "the world where water is XYZ" we are *misdescribing* something which we should rather call a "way our world might have turned out", but did not. *Given* that it did not, it could not have: "water is H_2O " is, *if* true, necessarily true. What we are imagining when we imagine water

*University of Geneva, Switzerland, philipp.keller@lettres.unige.ch.

¹Purely contingent or accidental truth of their doctrine, however, will not satisfy physicalists. There is a distinction to be made between worlds that count for or against physicalism and those that do not. But even then seems Chalmers illegitimately to infer (real) possibility from (mere) conceivability.

having another chemical structure than it actually has is not a possible world, a way our world could have been, but an impossible world, a way our world could have turned out but did not. Twin Earth, if it is to be possible, is not a world with water, but a world with a clear, drinkable liquid in its lakes and oceans, that we *would have called* “water” if we lived there. As we live here not there, however, the liquid we referred to when we actually fixed the reference of “water” had the chemical structure H_2O and so “water”, being a ‘rigid designator’, *actually* refers in all worlds to H_2O , though it could have done otherwise. For Kripke, the “might” in “water might not be H_2O ” is *merely* epistemic and represents “a present state of ignorance or uncertainty” (NN: 307, 103): given the evidence before the empirical investigation of the nature of water, it could have turned out that water is not H_2O - but in this case “water” would not mean what it actually means. The qualification “given the evidence” is hence crucial, for given *our* evidence, it is no longer true that water could have turned out not to be H_2O .²

David Chalmers has turned Kripke’s account upside down. Instead of a correct and an incorrect way of conceiving of (what one believes to be) a *p*-world, there are two intensions of *p*: the secondary intension of “water is H_2O ”, $[water\ is\ H_2O]_2$, is the set of possible worlds where H_2O is H_2O , the set of all worlds.³ Its primary intension, $[water\ is\ H_2O]_1$, is the set of worlds where the predominant clear drinkable liquid is composed of two hydrogen and one oxygen atom. Our world belongs to this latter set, but Twin Earth does not. In order to evaluate the primary intension of an expression α at a possible world, we have to “consider it as actual”: we have to figure out what we would refer to by α if that world turned out to be the actual world. To determine the primary intension of α , then, we have to extract from it by purely a priori means the “way in which [its] referent depends on how the actual world turns out” (CC: 12). To do this, we have to make use of what might be called a “presentation” of α , something giving us the way its reference depends on where we are. What might that be? Quite often, philosophers stay content with loose paraphrases.⁴ Let us abbreviate the definite description “the dominant clear, drinkable liquid...” by “watery stuff”. Instead of misdescribing Twin Earth as a world where water is not H_2O , we should rather call it a world where watery stuff has a different microphysical structure than it has on our world, thereby saying something which distinguishes our world from Twin Earth.

The illusion that “water is H_2O ” is contingent, is thus explained by Kripke and Chalmers in different ways: for Kripke, this is modal error, partially explained by the common, but mistaken, belief that necessary truths are a priori. This belief is based on the failure to distinguish metaphysical or logical and epistemic possibility and the failure to acknowledge the existence of *mere* epistemic possibility. Chalmers turns this the other way round. When we call “water is H_2O ” contingent, we are basically right, for its primary intension, expressed by “watery stuff is H_2O ” *is* contingent. Only its secondary intension, expressed by “ H_2O is H_2O ” is necessary.

²Kripke claims that “it might have turned out that *p*” entails “it could have been the case that *p*” (NN 141-142, 332).

³Following Chalmers’ usage in *The Conscious Mind*, I am speaking here and in the following of primary and secondary intensions of *linguistic items*. The notion of primary and secondary intensions of *concepts*, which Chalmers uses in *The Components of Contents*, seems derivative.

⁴Chalmers, e.g., says: “... it is not too far off the mark to say that the primary intension of my concept “water” picks out the dominant clear, drinkable liquid in the environment that stands in an appropriate causal relation to the center of a world, but this is far from perfect.” (CC: 11, cf. also CM: 57)

2 Chalmers' Argument against Physicalism

Chalmers' argument against what he calls 'type B materialism' proceeds from the following three premisses (MM: 4):

- (1) The physical facts do not a priori entail the phenomenal truths.
- (2) When S is a posteriori, S has a contingent primary intension.
- (3) If materialism is true, the physical facts necessitate all true propositions.

From (1) and (3) it follows that psychophysical identifications are necessary a posteriori truths. (2) and the conceivability of zombies entail that the primary intension of such an identification is not necessary. So there is a world where the primary intension of the physical antecedent is true while the primary intension of the phenomenal consequent is false. As primary and secondary intensions are evaluated over the same space of possible worlds, it follows that not all phenomenal facts are necessitated by physical facts and, by (3), materialism is false.

Chalmers' acknowledges that there is a loophole in the argument, as it has to be assumed that primary intensions of physical truths are still physical truths. If we were to deny this, however, we would get something like panpsychism which is not very materialistic either. I want to argue that there is a second loophole in the argument: how can we justify the step from the claim that a phenomenal truth Q is not a priori entailed by the physical truths to the claim that these physical truths do not necessitate the primary intension of Q ($[Q]_1$). It is based on a supposedly general principle, namely that Q and $[Q]_1$ are conceptually equivalent, in the sense that any necessitation of $[Q]_1$ gives us an a priori entailment of Q : if something necessitates $[Q]_1$, it a priori entails Q ; if P does not a priori entail Q , the conditional fails in a world considered as actual, and so $P \rightarrow [Q]_1$ fails in some possible world. This presupposes that we can conclude from

- (4) There is a P -world considered as actual where Q is false.

that

- (5) There is a P -world where $[Q]_1$ is false.

The step from (4) to (5) is the crucial step in the 'conceivability implies possibility' argument: if we can conceive (ideally, after enough rational reflection etc.) of a world falsifying q , there is a sense in which q is contingent.⁵

This crucial step is related to the illusion of contingency in the following way: To explain why an necessary truth may be falsely taken to be contingent, is to show that it is a posteriori in the sense of having a contingent primary intension. According to Chalmers, there are only illusions, but no outright hallucinations or delusions of possibility. In order for this to count as an explaining away of the illusion, we have to produce a conceptually equivalent statement which *is* contingent and to argue that it has been falsely taken for the original one. For theoretical identifications, stating that two rigid designators are coreferential,⁶ Chalmers suggests a general device for doing this: Whenever we have a necessary a posteriori truth $a = b$, its secondary

⁵I will use "conceivable" as dual to "a priori" and as "positively, ideally and primarily conceivable". I will ignore what Chalmers calls "secondary conceivability" (CP: 9) and come back to the question of idealization later.

⁶Like Kripke in the case of "pain" (NN 148, 337), Chalmers is committed to the claim that "consciousness" is a rigid designator, presumably referring to a natural kind. If it were not, then "consciousness is brain-state s " would be of same type than "watery stuff is H_2O ", i.e. be contingent a posteriori if true.

intension $[a]_2 = [b]_2$ is necessary, while its primary intension $[a]_1 = [b]_1$ is contingent. Whenever someone thinks $a = b$ is contingent, he really means that $[a]_1 = [b]_1$ which differs in modal properties, but is cognitively equivalent to the former in that $a = [a]_1$ and $b = [b]_1$ are *a priori*.

Recall that “watery stuff” has been introduced to account for the sense in which we seem to be able to imagine worlds where water is not H_2O . Kripke used it (or essentially equivalent variants thereof) to tell us that there really is *no sense* in which water might not have been H_2O : before learning that water actually *is* H_2O , we might have thought that “watery stuff is not H_2O ” is possible - but actually it is not. Our judgement to the contrary expressed a state of ignorance. Any qualitatively identical epistemic situation we are able to imagine *now* is one in which *our* watery stuff, i.e. water, is H_2O . Kripke’s argument for dualism rests on a disanalogy between “water” and “pain”, i.e. on the premiss that “pain” could not, in a qualitatively identical epistemic situation, turn out to refer to something which is not pain, i.e. which is unlike the phenomenal quality we experience when we are in pain. Kripke formulates this as a claim about the ways in which we, i.e. our linguistic community, pick out pain.⁷ This thesis is thus not entirely *a priori*, for we can imagine creatures that pick out (what is for them) pain in other ways. Kripke’s point is just that *we* cannot.

For Kripke, considering a possible world as actual is, in effect, imagining the meanings of our words to be different from what they actually are.⁸ If we know that p is a necessary *a posteriori* truth, “it might turn out that $\neg p$ ” may be true if p is not entailed *a priori* by our evidence for it, but “it might have turned out that $\neg p$ ” is false, given what we know, because we *did* conclude from our evidence that p (NN 143, n. 72, 354). The sense in which it is conceivable that water is not H_2O , then, is relative to times, persons and a certain body of evidence (NN: 35, 261). “Stick s is one meter long” is *a priori for* someone who used s to fix the reference of “one meter” (NN: 56, 274). IDENTITY AND NECESSITY 164

Chalmers, however, wants “watery stuff” to do more than this. He believes that it (or at least a relevantly similar refinement of it) captures the way in which “water is H_2O ” still is and ever will be a *posteriori*. In his view, the illusion of contingency consists in assimilating “water is H_2O ” to “watery stuff is H_2O ”. Kripke would have taken this to be a case of conceptual confusion, of ignorance about what our words actually mean.⁹ Chalmers, however, thinks that two sentences really are conceptually equivalent and differ only in that the first, but not the second, is metaphysically necessary. They are conceptually equivalent because they are connected by the following bridge principle, which is true (and known by us to be true) *a priori*:

(6) Water is watery stuff.

Chalmers’ two-dimensional argument against physicalism proceeds from the nonexistence of a functional specification ϕ of what it is to be in a phenomenal state ψ to the claim that the functional state ψ cannot be a *posteriori* identical to any brain state fulfilling the ϕ role. In doing this, he presupposes that ϕ and ψ are, if they pick out the same thing, *a priori* identical, thereby basing his argument on a *analogy* with the water case.

⁷“If any phenomenon is picked out in exactly the same way that we pick out pain, then that phenomenon *is* pain.” (NN: 153, 340)

⁸“... being put in a situation where we have exactly the same evidence, qualitatively speaking, it could have turned out that Hesperus was not Phosphorus; that is, in a counterfactual world in which ‘Hesperus’ and ‘Phosphorus’ were not used in the way that we use them, as names of this planet, but as names of some other objects, one could have had qualitatively identical evidence and concluded that ‘Hesperus’ and ‘Phosphorus’ named two different objects.” (NN: 104, 308)

⁹Polywater, however, allegedly discovered in the Soviet Union, is described by Kripke as water which is not watery (NN 129, 323).

3 A Dilemma for Chalmers

The dilemma I want to put up for Chalmers turns on the thesis that he is committed to the claim that primary intensions are - at least in principle - expressible. I have to arguments for this, the first *ad hominem* and the second programatic. If they were not he could not use the existence of an explanatory gap, i.e. the non-existence of a priori truths like (6) linking phenomenal and functional concepts, to argue against physicalism. The programatic argument is that apriority is tied to knowledge and knowledge to the *contents* of sentences. If we say of someone that he knows a priori that p , we do not merely ascribe to him knowledge that there is a sense in which p is and even cannot fail to be true, but we also ascribe knowledge of the truth actually expressed by p . If a priori knowledge of p were knowledge that p expresses a truth whatever it means, any epistemological interest for it would be lost.

Chalmers assimilates (6) to Kripke's cases of the contingent a priori, i.e. cases like "a is one meter long" where "a" denotes the standard meter in Paris, by pointing to which the reference of "one meter" in our language has been fixed. But whereas Kripke simply *stipulated* that he uses "a priori" in a way such that it applies to "a is ϕ " whenever ϕ is the description by which the reference of a has been fixed, Chalmers has a substantial notion of the a priori, one he wants to put to use in his argument against physicalism. I want to argue that no such substantial notion is to be had which makes (6) turn out a priori but contingent.

I want to challenge the general strategy of explaining away the illusion that " $a = b$ " is contingent by arguing that even in the water case, there is no sense of "a priori" in which it is both true that " $a = [a]_1$ " *is* and " $a = [a]_2$ " *is not* a priori. There is no X such that " X is H_2O " is a posteriori, i.e. conceivably false, and " X is water" is a priori, i.e. not conceivably false.

To explain why " $a = [a]_2$ " is *not* a priori, even given that we know it is true, we have to consider worlds as actual where we lack this knowledge, i.e. where we take something to be water that is not H_2O . To explain why " $a = [a]_1$ " *is* a priori, on the other hand, we have to restrict the realm of actual world candidates to those worlds where we *use the language as we do now*, i.e. where our words refer to sufficiently similar things than they do in our world. In order to do this, we have to exclude worlds where we are ignorant about what our words refer to: given that "water" uniquely refers to H_2O , we have to exclude worlds where we do not know that water is H_2O . For if we did not know that water is H_2O , we could consider a world as actual where water has quite different superficial properties than H_2O actually has, thereby falsifying the claim that it is a priori that water has the superficial properties it has.

In order not to misdescribe Twin Earth, Chalmers tell us, we have to say that it is a world where watery stuff is XYZ. But how can we be sure not to have misdescribed Twin Earth a second time, by using "watery stuff" instead of *its* primary intension? Chalmers' claim is not only that "[water]₁" and "watery stuff" are coreferential *in the actual world*. Instead, his idea is that "watery stuff" picks out *in every possible world* what is *there* the value of the primary intension of "water".¹⁰ For any world w , what water would be if w turned out to be actual is the watery stuff in w .¹¹ We may ask whether, in order to establish coreferentiality, we have to consider these worlds as actual or as counterfactual. Chalmers seems to think that it does not matter. I think it does.

I will try to show this by setting up a dilemma for Chalmers, corresponding to

¹⁰We are thus asked to understand "environment" in the quotation of the preceding footnote as a (universally bound) variable ranging over possible worlds.

¹¹"I will use "watery stuff" as a term of art to encapsulate the primary intension, whatever it is." (CM: 58) - But what does "encapsulate" mean?

two ways in which the term “reference-fixing description” may be used. In one way, it denotes the description, whatever it is, that has or might have been *used* to fix the reference of a particular term within our language community. In another way, it means whatever description is uniquely satisfied by the referent of the term. If we understand “watery stuff” as a reference-fixing description in the first sense, we take (6) to have the following form:

(7) Dthat(watery stuff) is watery stuff.

where “dthat” is *not* a rigidifying operator, but a directly referential singular term. This is the role Kaplan originally intended for “dthat” (Kaplan 1989: 579). The content of the accompanying description “the ϕ ” in “dthat(the ϕ)” is not part of the content of the singular term: the description merely serves as a “demonstrative surrogate” (1989: 581), completing the character and not the content of the term. It is, so to say, “off the record”, and not part of the semantics, but of the metasemantics of the referring term. I think this is the way Kripke intended his theory of the contingent a priori. I want to argue that it does not give us a substantial enough notion of a priori to do the work Chalmers wants it to do.

In the second sense of “watery stuff”, (6) comes down to:

(8) \mathcal{F} (watery stuff) is watery stuff.

where \mathcal{F} is a rigidifying operator (“fixedly”) that takes a description to form a singular term denoting in all possible worlds whatever uniquely satisfies the description. “ \mathcal{F} (the ϕ)” is, in Kaplan’s terms, “a rigid description which induces a complex ‘representation’ of the referent into the content” (1989: 580). We will have to ask ourselves what this representation is. I will claim that on no construal of “watery stuff”, however, (8) will come out interestingly a priori.

Here is another way to put the dilemma: In order to evaluate the truth of (6) in other possible worlds, we either consider these worlds as actual or as counterfactual. When we consider them as counterfactual, we are treating “watery stuff” as something like a demonstrative, a rigid designator referring to whatever has been pointed to when the reference of “water” was fixed. This is the dthat-construal of (6). Rigid designators refer in all possible worlds to the same object. So they are not epistemically transparent, for they distinguish between epistemic alternatives. If we consider these worlds as actual, however, we cannot, I will argue, pick out by “watery stuff” *its* primary intension and thus cannot capture the way in which we consider possible worlds ‘as actual’. In the first case, we have strong necessities, i.e. a posteriori necessities with metaphysically necessary primary intensions, at the heart of Chalmers’s theory. In the second, we get n -dimensionalism and lack any justification to stop at any particular n .

Let me make three preliminary remarks. My point is not just that primary intensions are ‘ineffable’ in natural languages in so far as that they are only capturable in terms of secondary intensions, a problem that might be met by introducing square brackets or similar devices stipulated to *uniquely* express the primary intension of their argument.¹² The deeper problem is that *any* way of specifying what the primary intension of a term is gives rise to the dilemma I want to put up for Chalmers in the following.¹³ The second caveat is that I have nothing against primary intensions in general: they work fine for indexicals, for example. What I am questioning is

¹²Cf. Chalmers: “The real narrow content is an abstract function from centered worlds into extensions, and is only imperfectly captured by a given description.” (CC: 38)

¹³Just a function, however, is not enough. Recall that Chalmers does not just present a semantical framework, but puts it to use against the physicalist, claiming that the physicalist cannot *specify* in functional terms what it is to be in a phenomenal state. This would not be a real challenge if

rather their supposedly intimate link to conceivability and the claim that they are scrutable by purely a priori means.¹⁴ Third, it is clear that, whatever else it is taken to mean, “watery stuff” cannot stand for a *purely* phenomenal concept, for then a world that only differs from our in that H₂O is not watery would already count as a zombie world, i.e. a world physically, but not phenomenally exactly like the actual world. “H₂O is watery”, however, is considered a posteriori on all sides.

I will now discuss in detail the two horns of my dilemma for Chalmers: primary intensions as descriptions by which reference has or could have been fixed and primary intensions as descriptions that are or can be rigidified.

4 The first horn: On the apriority of reference-fixing

Let us follow the first horn first and consider in what sense (6) may be called a priori if “watery stuff” is taken to be the description possibly *used* to fix the reference of “water”.

Like “water is H₂O”,¹⁵ “water is watery stuff” allows readings under which it is false. Stalnaker and Block (CDE: 10) have described cases in which “watery stuff” could fail to pick out a unique referent, due to nonphysical ghost water. We do not even have to consider far fetched possibilities, however. There is a chemical substance, heavy water or deuterium oxide, which is watery stuff, but not water, for it has the microstructure D₂O. It is odour- and tasteless and you could not tell it from water.¹⁶ So “water is watery stuff” is true only if “watery stuff” is taken to include a descriptive component excluding D₂O. As D₂O has been synthesized for the first time in 1933 and water has been discovered to be H₂O already in 1784, the description by which the reference of “water” has been fixed presumably did cover heavy water as well.

To this it might be replied that the actual reference-fixing description does not matter. What is important, however, is the disanalogy between a description mentioning only superficial properties that might have been used to fix the reference of “water” and a description of its microphysical properties. So let us now consider

the physicalist could come up just with a function, picking out in any possible world what has the functionality of the relevant phenomenal state there. So I do not agree with the claim that “the PI framework doesn’t need to take a stand on whether the PI can be captured by a description” (Chalmers, <http://www.u.arizona.edu/~chalmers/class/596b/week4.txt>). If primary intensions are not specifiable at all, then how could we have cognitive access to them?

¹⁴Cf. Chalmers: “The primary intension of a concept, unlike the secondary intension, is independent of empirical factors: the intension *specifies* how reference depends on the way the external world turns out, so it does not itself depend on the way the external world turns out.” (CM: 57)

¹⁵As Avrum Stroll has reminded us (1989), if ice is in the same way ‘identical to’ (i.e. composed of) H₂O than water, then water is ice. But even if ice and steam count, under a relaxed sense, as water, water is *not only* H₂O, but contains various minerals, mud etc. So only *pure* water is H₂O. Pure water, however, is not what we most often mean by “water”. If I order a glass of water in a restaurant, I would not like to get pure water, because pure water is not healthy to drink. So “water is H₂O” is true only if it is taken to mean “H₂O is the dominant chemical substance water is composed of”. “Dominant”, however is vague. Consider the sentence “air is O₂”. There are many uses of “air” in which “air” is taken to refer to oxygen. Take, e.g., “The people in the submarine died because there was not enough air left for them to breathe.” Nevertheless, air consists only to 21% of oxygen and it seems false to say that science discovered that air *is* oxygen. Moreover, air and oxygen share some but not all of their macrophysical properties. Like air, oxygen is odourless, transparent and breathable, but smoking a cigarette in a room filled with pure oxygen is even more unhealthy than smoking it in a room filled with air. It seems that on some uses of “air”, “air is O₂” is necessary, but on others it is not.

¹⁶Apparently, mice can tell them apart. This suits them well, for D₂O makes them infertile. We also know that small fish cannot live very well in heavy water.

whether we find a disanalogy between “water is H_2O ”, supposedly necessary but a posteriori, and “water is watery stuff”, supposedly contingent but a priori.

What does it mean to imagine discovering that water is something other than H_2O ? In what sense is it conceivable that water is not H_2O ? We have to distinguish different cases: First, I take it that you cannot imagine of water and of H_2O that they are distinct, for this would mean to imagine of one and the same thing that it differs from itself. You can, second, imagine of water that it has a different molecular structure than it actually has. In this case, you are treating “ H_2O ” as a description. This can be seen from the fact that it will not, under this reading, be necessary that what actually *is* H_2O , i.e. water, is H_2O , i.e. has the microphysical structure it has. You can also imagine of water that it has another name than it actually has and, if you keep the language of chemistry ‘fixed’ while doing this, you thereby imagine that it has a different molecular structure. Third, you can imagine of H_2O that it has different macrophysical properties. For this to count as a case of imagining of H_2O that it is not water, you have to treat “water” as a description of those macrophysical properties you are imagining H_2O to lack. You can equally imagine of H_2O that it might have been called other than “water”. If you keep your language ‘fixed’, you are thereby imagining that the description by which the reference of “water” has been fixed is true of something other than water.

The upshot of this is that in no case it will be the same thing that is both necessary and conceivably false.¹⁷ The obvious reply is that we have, in order to evaluate the claim that “water is H_2O ” is informative, to *bracket* our knowledge that it is true. But what does this mean? It clearly will not do to imagine a situation where I know everything about the microphysical structure of water except that it is *called* H_2O . So I have to imagine being ignorant either about water or about its microphysical structure. If water *is* H_2O , this amounts to the same thing. To explain the possible truth of “I do not know that water is H_2O ” is thus just to explain the possible truth of “I do not know what water is”, and to explain this we have, given that water is H_2O in all possible worlds, to solve the problem of logical omniscience.¹⁸

Let us now take the dual, the supposedly contingent but a priori truth “water is watery stuff”, where “watery stuff” abbreviates a description uniquely true of water, but not including its microphysical properties. It seems to me that we get exactly the same picture than before. You cannot, first, imagine of water and of water stuff that they are distinct. You can, however, imagine of water, i.e. of H_2O , that it has different macrophysical properties than it actually has. You can equally imagine of watery stuff, i.e. of water, that it has different microphysical properties than it actually has and thereby, treating “water” as a description of these microphysical properties, imagine that watery stuff is not water. There simply does not seem to be a reason to privilege the macro- over the microphysical description, taking one to be somehow part of the meaning of “water” in a sense that the other is not.

Something must have gone wrong. The diagnosis, of course, is that “watery stuff” is supposed to have played a special role in the fixation of the reference of “water” that distinguishes it from “ H_2O ” which has been coined long after the inclusion of the word “water” into our language. The claim, then, is that (8) is a case involving

¹⁷I take this to be compatible with Kripke’s insistence, in *Naming and Necessity*, on the importance to distinguish between metaphysics and epistemology, between what is necessary and what is a priori. The mistake of those who try to refute the necessity of “water is H_2O ” by imagining a world where watery stuff has a different microphysical structure is that they do not imagine a world where what Kripke claims to be necessary is false. In this sense, they are misdescribing the world they imagine, if they take it to be a counterexample to the necessity thesis.

¹⁸This problem is even particularly pressing in the present context. For if “ $a = b$ ” is true, then it is a theorem and so you know it (as it is true in all your epistemic alternatives). So if you really are ignorant about it, it is false. Making this reasoning, you are no longer ignorant about it.

some sort of a “descriptive name” in the sense of Evans. Evans invited us to stipulate that “Julius” denotes the inventor of the zip. He then contends that

(9) Julius invented the zip.

is a priori and what he calls ‘superficially contingent’. But there are two problems here. The first, recognized by Evans, is that there might have been no or no unique inventor of the zip. So what is a priori, if anything, is not (9) but “If there was a unique inventor of the zip, then Julius invented the zip”. But there is a second problem too. Let us agree that (7) has the same logical form than (9), namely:

(10) $\phi(\text{dthat}(\text{the } \phi))$

and let us assume that “the ϕ ” picks out a unique person in our world. The problem is that we have no assurance whatsoever that this person invented the zip, i.e. that (10) is true.¹⁹ Keith Donnellan is famous for having recognized that, in order to refer to something by a definite description, this description does not have to be true of it. Donnellan calls such uses “referential” and they have been taken to be the paradigm cases of the use of descriptions to pick out, rather than to describe something. We do not have to know anything about the inventor of the zip to *stipulate* that he is to be called “Julius”. If we understand “the inventor of the zip” in “Let “Julius” denote the inventor of the zip” in the *referential* and not in the attributive sense, we cannot be sure that we named thereby the actual inventor of the zip, given that we do not know who the inventor of the zip is.

If we, however, *stipulate* that our reference-fixing description be true of the thing reference to which is fixed, we get an empty notion of the a priori. Let me tell you why. The original argument for the apriority of (9) had the following form:

(11) Let “Julius” denote whoever invented the zip. Then it is a priori that Julius
invented the zip.

When we wonder whether (9) is a priori, we have to ask ourselves whether we can consider a world w as actual which is such that there is someone x called “Julius” in w and it is not true of x in w that he or she invented the zip. Whether or not it is a priori that Julius invented the zip, thus depends on whether or not it is a priori that “Julius” denotes the inventor of the zip, that is on whether or not we can conceive of a world as actual where “Julius” does not denote the inventor of the zip. In one sense, this is easy. It is perfectly possible to conceive of a world as actual where the person who invented the zip is not called “Julius” (the actual world, e.g., will do).

But we *stipulated* “Julius” to be a name of the inventor of the zip, whoever it was who invented the zip. The crucial question now is: what does this stipulation amount to? It can mean two different things. Either it amounts to excluding any worlds from our consideration where the person who invented the zip there is not called “Julius” there. This, however, cannot be the intended interpretation, for it is easy to conceive of worlds as actual where somebody other than the inventor of the zip is called “Julius”. Instead, the stipulation amounts to excluding any worlds where the person *we would have called* “Julius” if we lived there did not invent the zip there. Here we have a double modality. There are two ways to cash it out. Either we take the stipulation to fix a *two-dimensional* intension for “Julius”. This would mean that we thereby fix not only our actual language, but also the language we would speak in

¹⁹This is especially clear in cases of semantic deference: the primary intension of “elm”, according to Chalmers, picks out whatever trees those around the speaker refer to as “elms”. It is clear that I can be wrong about what they *actually* refer to and so not be referring to what I intend to refer to.

counterfactual circumstances. Alternatively, we suppose that the stipulation is made in any world which we consider as actual. What it means to “keep the language fixed”, then, is to stipulate that they use the same words with the same meanings to fix the reference of “Julius” than we do.²⁰ To keep the language fixed, then, is to exclude worlds w where “Julius” does denote somebody other than whoever invented the zip in w . Whether we fix the two-dimensional intension of “Julius” or whether we iterate the naming procedure in each world we consider as actual, what may properly be called “a priori” is not (9) but the following conditional:

(12) Let “Julius” denote the inventor of the zip. Then Julius invented the zip.

(12), in contrast to (9), however, is not only a priori but also necessary. There is no world, conceived as actual or not, where there is a unique inventor of the zip, “Julius” is used as a name for him and “Julius invented the zip” is false. So (9) comes out a priori under an empty notion of “a priori”: something is a priori in this sense in virtue of being of the form “if ‘p’ is true, then p”.

So we are left with a choice: either we *presuppose* that we would have called “Julius” some other person if some other person would have invented the zip, thereby making (9) *trivially* a priori or we allow for the possibility that we may falsely believe that we have named a but in fact have named b . To know whether a world w is a world where “Julius invented the zip” is false, we would then have to know whether a invented the zip. The upshot of this, I think, is that it does not follow from “if a were not ϕ , we would not call it ‘a’” that $\phi(a)$ is a priori.²¹

The point carries over to “water is watery stuff”. Either you allow for the possibility that the reference of water has been fixed by a description which is not in fact true of it. If you stipulate, on the other hand, that “water” denotes watery stuff whatever it is, you do not only keep your language fixed, but a lot of empirical knowledge as well, namely your *de re* knowledge of water that it is watery. So you are no longer considering worlds as “actual”, but you are considering them as counterfactual, as ways the actual world might be *given* that water is watery. You are then supposing that the primary intension of “water” is given by the *secondary* intension of “watery stuff”, i.e. that $[\text{watery stuff}]_2 = [\text{water}]_1$. This makes “watery stuff” behave very much like a rigid designator. In every possible world, an utterance of “water” would thus refer to what is there the predominant liquid in the lakes and the oceans. This cannot be right, however, for it is easy to imagine worlds where water has become sparse, where only small samples of it are kept in laboratories and the lakes are filled with, say, petrol, but where the word “water” still applies to what it applies here, namely to H_2O . Entertaining the hypothesis that water might become sparse is not envisaging a linguistic change.

It is, e.g., possible to imagine a world where the clear drinkable liquid in the lakes and oceans is not what we (or members of our linguistic community or experts or whoever) pointed to when we fixed the reference of “water”. It is possible that,

²⁰This observation is Stalnaker’s: “... for this representation [of (9) as a contingent a priori truth] to be right, we must add that the stipulation in question was made in each of the worlds i , j and k . One who did not know about the stipulation, or did not understand it, would not know that the statement was true.” (1999: 15)

²¹The reason for this is that we use proper names for all sorts of purposes. If I say to you “if I were Maradona, I would marry Britney Spears” and I believe that Maradona is the best football player in the world, it may well be that if I knew that Schmaradona, but not Maradona, is the best football player in the world, I would instead say “if I were Schmaradona, I would marry Britney Spears”. But from this it does not follow that it is a priori that Maradona is the best football player in the world. Note also that the converse of the above principle is false as well. Take, e.g., the supposedly contingent a priori truth “I am the speaker”. If I say to you in a conversation “if I were you, I would do p ” I am not misdescribing a possibility that should be properly described by “if the speaker were the audience, the speaker would do p ”. Thanks to Andri Töndury for the hint.

unbeknownst to us, the reference of “water” has been fixed in other ways than we think it is or was. It is tempting to reply that it is (merely) *linguistic* knowledge that we would be lacking in this case. There are two replies to this: First, linguistic knowledge is not in an interesting sense a priori; second, if it were a priori, it would give us a priori empirical knowledge. ‘Linguistic’ knowledge in this use of the word, is knowledge about how reference is fixed, where ‘fixing the reference’ is an *activity* of certain members of my linguistic community, an empirical fact which I can easily imagine to have been otherwise. Kripke, Donnellan and others have described and imagined cases where originally descriptive names have become proper. How do we realize this? It certainly depends on whether a major or the most important part of our linguistic community is prepared to accept sentences of the form “*a* is not ϕ ” (where “the ϕ ” is the description by which the reference of “*a*” has been fixed) as possibly true. The converse, however, is not true: even if they regard such sentences as plainly false, this does not show that they take it to be of the form “the ϕ is not ϕ ”. The second point pertains to a retreat position of the advocate of the contingent a priori sketched by Donnellan (1979: 55). Cannot we at least say that the following is a priori?

(13) Provided the ϕ exists, “*a* is the ϕ ” expresses a contingent truth.

(13) is supposed to speak about a particular language *L* and more specifically about the truth-theory for that language. We know a priori that any truth-theory must give us all instances of the *T* scheme and so we know, if *L* is our own language, that (13 is true, provided that the ϕ exists, only if *a* is the ϕ .

To this is may be replied that “watery stuff” contains an indexical element, that it denotes the clear drinkable liquid in *my*, or *the speaker’s* environment.²² Either the relevant features of the speaker’s environment can be specified in non-indexical language or they cannot. If they can, as Chalmers seems to think (MM: 5), nothing crucial changes. If they cannot, “watery stuff” distinguishes between epistemic alternatives,²³ so it is not what we could call “epistemically transparent”. It was precisely their epistemic transparency, however, that led Chalmers to introduce primary intensions in the first place. If “watery stuff” is not epistemically transparent, if it is possible that you are unable to tell whether its application conditions are fulfilled, “water is watery stuff” is a posteriori, while it has a metaphysically necessary primary intension.

We can, of course, as Kripke did, *stipulate* to use “a priori” as a predicate that applies to all and only the sentences of the form *Fa*, where “the *F*” is the description whereby the reference *a* has been fixed for all, many or most of us (NN: 78, 290). But there are good reasons not to accept such a stipulation. For we also use “a priori” to speak about certain kinds of knowledge and this use is undercut by the stipulation. If “the man who discovered the Peano axioms” fixes the reference of “Peano”, some of us both know a priori that Peano is the man who discovered the Peano axioms *and* know a posteriori that Dedekind discovered them. So they know that Peano is Dedekind, which is false. Similarly, if we *know* a priori that Gödel discovered the incompleteness of arithmetic, then we know that arithmetic is incomplete. If a mistake were found in the proof, this might turn out to be false.

²²This is suggested by Chalmers’ remark (CC: 5) that we have to evaluate the primary intension at *centered* worlds, worlds with an index fixing the speaker, the time and location of the utterance in question etc. We are thus supposed to include into the ‘center’ a sample of water, e.g. the one that members of our linguistic community pointed to when they introduced “water” into our language.

²³Two worlds are epistemic alternatives if they cannot be told apart by their inhabitants, if there is nothing with respect to which their inhabitants could distinguish them.

5 The second horn: On the rigidity of meaning this

Let us now try the second horn, where we take “water is watery stuff” to be of the form (8). Unlike (10), “ $\phi(\mathcal{F}(\text{the } \phi))$ ” (“whatever is ϕ is ϕ ”) is a truth of logic and can thus plausibly be taken to be a priori. The problem now is not with the a priori, but with contingency: how could “ $\phi(\mathcal{F}(\text{the } \phi))$ ” fail to be true? Where does the disanalogy between “watery stuff($\mathcal{F}(\text{watery stuff})$)” and “ $\text{H}_2\text{O}(\mathcal{F}(\text{H}_2\text{O}))$ ” come from?

The obvious response to this question brings in the two-dimensional framework. In order to imagine a world where water is not H_2O , you have to consider it as actual and not as counterfactual, whereas to imagine a world where water is not watery stuff you have to consider it as counterfactual and you cannot consider it as actual. What does this contrast amount to?

In Kripke’s framework, it amounts, given that you know that water is H_2O , to the difference between imagining having had another past and imagining being bound to have another future than you think you will have. While considering a world where water is not watery stuff as actual, you are imagining that “water” meant something other than what you took it to mean all along. You then ask yourself what water would turn out to be in such a world. While it could not turn out not to be watery, it could well turn out not to be H_2O . While considering a world where water is not H_2O as counterfactual, however, you imagine that what we call “water” is not composed of two hydrogen and one oxygen atom. Given that you know that it is, you cannot imagine that. You can, however, imagine finding out that you did not really *know* that the chemical substance you call “water” has the macrophysical properties you believe it has, that is you can imagine finding out that H_2O is not watery.

In Chalmers’ framework, the difference between “watery stuff($\mathcal{F}(\text{watery stuff})$)” and “ $\text{H}_2\text{O}(\mathcal{F}(\text{H}_2\text{O}))$ ” is that the definite descriptions thus play different roles in the fixation of the reference of “water”: you use “watery stuff” to rigidify its primary, and “ H_2O ” to rigidify its secondary intension. So we get the following variant of (8):

$$(14) \quad \dagger(\text{watery stuff}) \text{ is watery stuff.}$$

where \dagger is a function mapping a singular term a to another singular term $\dagger a$ denoting in every possible world the semantic value a would have if used there, i.e. a function projecting its diagonal intension onto the horizontal.

For “water is H_2O ”, we get

$$(15) \quad \ddagger(\text{H}_2\text{O}) \text{ is } \text{H}_2\text{O.}$$

where \ddagger is Stalnaker’s ‘upside-down dagger’, i.e. a function mapping a singular term a to another singular term $\ddagger a$ denoting in every possible world the semantic value of a attributed to your actual use of the term, i.e. a function projecting the diagonal proposition onto the vertical (Stalnaker 1968: 83n.).

Given that “water” has the primary and secondary intensions it is claimed to have, (14) and (15) are true in the actual world. But Chalmers claims more than this: he claims that (14) and (15) are not only true but a priori and a posteriori respectively. To be a priori, it is claimed, is to have the same primary intension. But what are the primary intensions of “watery stuff” and “ H_2O ”?

Let there be three worlds, the actual world, where water is watery stuff and H_2O , Twin Earth, where water the watery stuff is XYZ, and a third world, where the watery stuff is ABC. So “water” has the following two-dimensional intension:

H_2O	H_2O	H_2O
XYZ	XYZ	XYZ
ABC	ABC	ABC

“Watery stuff” is stipulated to express the diagonal intension of “water” in the actual world. If uttered in Twin Earth, however, it would express the diagonal intension of their concept of “water” and so it would in the third world. I use subscripts to indicate the secondary intension of the concept Twin-Earthian use to express the primary intension of their watery concept (that we may express by “twatery stuff”). So we get the following matrix:

H ₂ O	XYZ	ABC
H ₂ O ₁	XYZ ₁	XYZ ₁
H ₂ O ₂	ABC ₂	ABC ₂

In the two-dimensional framework developed by Davies and Humberstone, this amounts to the following:

$$(16) \mathcal{F}A(\text{Water is watery stuff})$$

where $\mathcal{F}A$ (“fixedly actual”) is defined as $\models_y^x \mathcal{F}A\alpha :\Leftrightarrow \forall y \models_y^y \alpha$ (where “ $\models_y^x \alpha$ ” means “ α , as evaluated in y is true of x ”). $\mathcal{F}A$ corresponds to $\Box\ddagger$, where \ddagger is a function mapping the diagonal proposition of a propositional concept onto the horizontal.

$$(17) \mathcal{F}\Box(\text{Water is H}_2\text{O})$$

where $\mathcal{F}\Box$ is defined as $\models_y^x \mathcal{F}\Box\alpha :\Leftrightarrow \forall w, z \models_z^w \alpha$. $\mathcal{F}\Box$ (or, equivalently, $\Box\mathcal{F}$) corresponds to $\Box\ddagger$, where \ddagger is a function mapping the diagonal proposition of a propositional concept onto the vertical.

In a more uniform notion, this amounts to the following:

$$(18) \Box(\ddagger\text{Water} = \ddagger\ddagger \text{watery stuff})$$

$$(19) \Box(\ddagger\text{Water} = \ddagger\ddagger \text{H}_2\text{O})$$

Does these sentences express the claim that (14) and (15) are respectively a priori and a posteriori? I think not. For (18) and (19) are true iff (14) and (15) are true, respectively, and, both projection operators being idempotent, they are even true iff “watery stuff is watery stuff” and “H₂O is H₂O” are true.

This, however, means that “water is watery stuff” is a priori if it is true.²⁴ If we are ignorant about which of various actual world candidates is the real actual world, we cannot specify one world as actual and so the two-dimensional analysis does not even get from the ground.

Here is another way to make the same point:

The problem, however, is that this is not enough: we may well conceive of worlds where (16) is false, i.e. worlds where H₂O is not watery. In order to make (16) necessary, we have to restrict ourselves to worlds where “water” denotes what is there the watery stuff. In order to evaluate “water is watery stuff” at a possible world considered as actual, we have to hold the meaning of “water” constant. In Chalmers’ words, we have to retain the concept from the “real actual world” (CC: 8, 42). In order to do this, we have to know what the relevant meaning of “water” is. But plainly, we simply do not know which of the different actual world candidates *is* the real actual world. But does it matter? Chalmers thinks it does not,²⁵ and I think it does. Here is why: A statement p is conceivable (i.e. $\neg p$ is not a priori), Chalmers claims, iff it’s primary intension $[p]_1$ is possible. The primary intension is supposed to represent the a priori accessible content of p . What we are doing, then,

²⁴In Davies’ and Humberstone’s two-dimensional logic, we cannot conclude α from $\mathcal{F}A\alpha$. All we can conclude from it is $A\alpha$, which is true iff α is true in the actual world, whatever this is.

²⁵

when we conceive of p , is conceiving of a world where some other phrase q we take to express the primary intension of p is true. We are conceiving of the truth not of p but of the conceptually equivalent “[p]₁”. The “conceivability implies possibility” thesis then amounts to:

(20) “[p]₁” is conceivable \iff “[p]₁” is possible

By the definition of primary intensions, conceiving of “[p]₁” is conceiving of something that expresses *its* primary intension, “[[p]₁]₁”. So (20) gives us:

(21) “[p]₁” is possible \iff “[[p]₁]₁” is possible

Though Chalmers is not committed to the claim that [p]₁ = [[p]₁]₁,²⁶ he is, I think, committed to (21).

On the two-dimensional analysis of what it is to conceive that p , one only has to believe that p expresses some truth or other at some world. One has to imagine that the sentence, whatever it means, expresses a truth. For this to count as a way of imagining the truth of what is expressed by the sentence, however, one has to know what the sentence means. And this, if Kripke is right, is a posteriori knowledge, depending on how the real actual world turns out.

In order to know that our evidence for the possibility of p does not just depend on our ignorance, we have to know that what [p]₁ expresses is possible. To know this, we have to conceive of a world where it is true. For we have no other access to what is possible than via conceivability. To claim otherwise would be to claim that brute facts about the way we speak constrain the realm of possibilities. To conceive of a world where [p]₁ is true, however, we have to presuppose that what [[p]₁]₁ expresses is possible. So we get a regress.

Suppose, then, that “watery stuff” picks out the primary intension of “water” in every world *considered as actual*, i.e. that [watery stuff]₁ = [water]₁.²⁷ The primary intension specifies how the concept expressed depends on the world. This, however, is empirical knowledge: whether or not the the meaning of “watery stuff” depends on the world *in the same way* than the meaning of “water” is a fact we know in virtue of being member of a certain linguistic community, of having learned the language we speak in one way but not in another.

To see the problem, consider the following sentences:

(22) If XYZ is watery, then XYZ is water.

(23) If H₂O is not watery, then water is not watery.

Chalmers argues from the truth of (22) to the aposterioricity of “water is H₂O” and from there to the claim that the primary intension of “Water is not H₂O” is non-empty. To evaluate the truth of (22), he uses the Ramsey test: accept hypothetically the antecedent and check whether you can conclude to the consequent (CP: 13). So if “water is watery stuff” is a priori, (23) had better be false. But is it? I am not sure. Assume hypothetically that the antecedent of (23) is true. Given that you know that

²⁶With reference to Kripke, Chalmers claims that the primary and secondary intensions of “consciousness” and of phenomenal concepts in general coincide : “What it takes for a state to be a conscious experience in the actual world is for it to have a phenomenal feel, and what it takes for something to be a conscious experience in a counterfactual world is for it to have a phenomenal feel.” (CM 133; cf. also MMM: 6).

²⁷Chalmers probably has this in mind when he emphasizes that the primary intension is independent of empirical factors: “The primary intension of a concept, unlike the secondary intension, is independent of empirical factors: the intension *specifies* how reference depends on the way the external world turns out, so it does not itself depend on the way the external world turns out.” (CM: 57)

water *is* H₂O, cannot you therefrom conclude that water is not watery? But what about Twin Earth, where H₂O is not watery, but what you would call “water” there, namely XYZ, is? Twin Earth is a world where you do not know that water is H₂O and where you therefore consider it possible that water is not H₂O, but say, XYZ.

The problem thus is the following: in order to evaluate primary intensions at worlds considered as actual, we cannot but evaluate *what we believe are the primary intensions of our words*. In this belief, however, we can be wrong.

Chalmers provides the following argument for the apriority of our knowledge of primary intensions:

“Given that we have the ability to know what our concepts refer to when we know how the actual world turns out, then we have the ability to know what our concepts would refer to *if* the actual world turned out in various ways.” (CM 59-60)

This does not seem right, however. It may be that, given that we know how the actual world turns out (e.g. that water is H₂O), we know what our words refer to (e.g. “water” to H₂O). It does not follow, however, that we have the further knowledge that, for any possible way the actual world turns out, our words would refer to these things. Linguistic knowledge is knowledge about one particular this-worldly language, not knowledge about a whole range of different and at the same time sufficiently similar languages.²⁸

When we falsely take a primary intension to be nonempty we take a world verifying it to be an actual world candidate which is not such a candidate. This could be so because “being an actual world candidate” is a non-empty condition on worlds. This is, e.g., the case in Kaplan’s *Logic of Demonstratives* where contexts of utterance are required to contain a speaker in order to make “I exist” come out as logically true. As Chalmers takes primary and secondary intensions to be total functions defined over the same space of possible worlds, I will not pursue this line further. Another reason for being wrong about a primary intension, however, is that we may falsely believe that something is an actual world candidate which is not even possible, but belongs to the so-called “outer sphere” of possibility.

Chalmers agrees that we have, when evaluation apriority claims, to make sure that we are not misdescribing worlds we consider as actual.²⁹ But how can we do this?

“The fact that p might, for all I know, be true in the actual world, is just irrelevant to the issue whether it is true in some possible world.” (Yablo 1993: 9)

To rule out linguistic error, Chalmers would need an extremely strong notion of the a priori according to which p is a priori iff it is not possible to believe $\neg p$. The corresponding notion of conceivability as possible believability, however, is far

²⁸Note that the point is not only that you cannot infer $K(\phi \rightarrow \psi)$ from $K(\phi) \rightarrow K(\psi)$, but that you cannot infer $\forall F K(\phi(F) \rightarrow \psi(F))$ from $\exists G (K(\phi(G)) \rightarrow K(\psi(G)))$. It may seem that I am advocating a quite general meaning scepticism in my claim that we can falsely take a primary intension to be non-empty. If the meaning in question is taken to be that of “water”, e.g., then it is true that I advocate a moderate scepticism in the way of Kripke. But everybody does, so this does not count against my point. If it is that of “watery stuff”, however, I indeed think that we have a clear grasp of its primary intension. As “watery stuff” is a technical philosophical term, however, this claim is in no way contrary to what anyone, excluding philosophers, believes.

²⁹“On this view of conceivability, the conceivability of a statement involves two things: first, the conceivability of a relevant world, and second, the truth of the statement in that world. It follows that in making conceivability judgements, one has to make sure that one describes the world that one is conceiving correctly, by properly evaluating the truth of a statement in the world. One might at first glance think is it conceivable that Goldback’s conjecture is false, by conceiving of a world where mathematicians announce it to be so; but if in fact Goldback’s conjecture is true, then one is *misdescribing* this world; it is really a world in which the conjecture is true and some mathematicians make a mistake.” (CM: 68)

too weak to imply possibility, for many impossibilities may and have been believed. That we find ourselves unable to rule out a possibility does not show that it is a possibility we think we cannot rule out rather than an impossibility.

Does idealization help? Chalmers claims that ideal positive and primary conceivability of p , that is conceivability after ideal rational reflection of a scenario that would verify p if actual, entails that p is possible. So it could be replied that the regress only arises in the case of less than ideal conceivability. I do not think that is true. To see why, consider Yablo's analogy of conceiving with perceiving:

“Just as to perceive that p is to be in a state that (i) is veridical only if p , and that (ii) moves you to believe that p , to find p conceivable is to be in a state which (i) is veridical only if possibly p , and (ii) moves you to believe that p is possible.” (1993: 7)

There is no a priori guarantee that perceptual error, even in principle and under ideal perceptual circumstances, can be ruled out. Why should there be such a guarantee in the case of conceivability? But the case is worse than this: with regard to perception, we are not forced to sceptical conclusions from the absence of such an a priori guarantee: even if there is no guarantee that p is true, given that I perceive that p under ideal circumstances, it may still be and probably is the case that my perceiving that p gives strong, but defeasible, evidence that p . Chalmers, however, cannot retreat to the claim that conceivability (even conceivability under ideal conditions) is defeasible *evidence* for possibility. His analysis of conceivability in terms of primary intensions entails that in any case where I really conceive that p , the primary intension of p is non-empty. So Chalmers is committed to the claim that under certain circumstances, conceivability is an *infallible* guide to possibility.

It seems that we are left with two options, corresponding to the two horns of my dilemma. Either we concede that (6) is not a priori after all, in virtue of the fact we can conceive of a world as actual where it is false, and thus concede that aprioricity does not imply contingency of primary intensions. Or we stipulate that primary intensions capture that cognitive content of statements and allow for a regress, corresponding to the non-transitivity of aprioricity. In the last part of my talk I would like to explore both these options a little bit further, beginning with the second first.

6 n -Dimensionalism

Suppose watery stuff in w is not H_2O , but XYZ, but that the inhabitants of w nevertheless use “water” to refer to H_2O . They certainly can consider a world v as actual (e.g. their own actual world w) where H_2O is not watery, in the same way as we can imagine a world where H_2O is not watery. So v is a world, considered as actual from w , where the substance inhabitants of w call “water” is not watery. It is not, however, a world where what *we* call “water” is not watery, for what we would call “water” in w is not H_2O , but XYZ. The inhabitants of w differ from us in their concept of water, not because they mean by it something other than H_2O (for they do not), but because they use it to refer to something which is not watery *by our standards*, i.e. which does not fill their lakes etc. Note, however, that this difference only shows up if we *presuppose* that it is H_2O and not XYZ that fills *our* lakes etc. If we were wrong about this, we could, for all we know, be inhabitants of w .

Even if we are not wrong about H_2O 's filling our lakes etc., we *can* easily *imagine* ourselves to be, i.e. we can conceive of w as actual. We can consider a world as actual in which we *falsely* believe that watery stuff is H_2O and which is thus a world in which we falsely believe to be able to distinguish between what we then would take

to be our actual world and w . If we conceive of a world as actual in which we are, but do not think we are, using “water” to refer to XYZ and where we imagine that XYZ is not watery, we conceive of a world as actual where “water is watery stuff” is not a priori. So it is not a priori that “water is watery stuff” is a priori.

It will be replied that we, as in the case of a posteriori truths, we have to restrict our claim to a conditional. “Water is H₂O” is necessary, *if* it is true. Can we thus say that “water is watery stuff” is a priori *if* it is true? Doing this we would trivialize our notion of “a priori”. Note that there is a crucial asymmetry between the two cases:

(24) If p is true, p is necessarily true.

(25) If p is true, p is a priori.

Kripke argues that we know the truth of (24) a priori, by philosophical analysis (NN 109, 311). But what about (25)? Kripke admits that even when we have fixed the reference of F with respect to a sample, it is still an empirical question whether all or most of the items in the sample are F .³⁰ What we have to presuppose in the case of “water is watery stuff” is not only that it is true but that it is *known* to be true: aprioricity does not, like necessity, qualify truth but knowledge.³¹ *Given* that we know that it is true, it may be called “a priori”, but vacuously so, for every p is entailed by a body of evidence which includes p itself.

Imagining this, we distinguish the primary from the secondary intension of “watery stuff”. So we get a regress, in many respects similar to the regress of Fregean senses. This has been remarked by Davies and Humberstone³² and is familiar from the literature on Frege.³³

The analysis of conceivability in terms of non-empty primary intensions seems to force upon us a non-S5 logic of metaphysical modality. For we conceive of a world as actual where water is not H₂O, so there is a world where watery stuff is not H₂O. We can equally conceive of a world as actual where science tells us that water is necessarily H₂O. So there is a world where watery stuff is necessarily H₂O. So $\Diamond\Box\alpha$ does not imply $\Box\alpha$. It has to be, however, transitive, at least for necessary a posteriori truths like “water is H₂O”, for it is supposed to accommodate for Kripke’s observation that such sentences are, if true, necessarily true. So in any world w I consider as actual, if “water is H₂O”, as understood by us, is true in w , then it is true in all worlds reachable from w . If it is possible that water is *possibly* H₂O, then science has not found out what it claims to have found out.

7 This-worldly necessity

Closely associated with Chalmers’ “conceivability implies possibility” claim is a doctrine he calls “modal rationalism”. ‘Modal rationalism’, as I understand the term, denies that the space of metaphysically possible worlds could be a proper subset of the space of conceptually possible worlds, i.e. worlds that might (ideally, after reflection etc.) be *taken to be possible*.³⁴

³⁰“A priori, all we can say is that it is an empirical matter whether the characteristics originally associated with the kind apply to its members universally, or even ever, and whether they are in fact jointly sufficient for membership in the kind.” (NN 137, 329)

³¹If p is true, but not known to be true, how could it then be a priori (on any reasonable sense of “a priori”)?

³²“For the record, we should remark that the process which yielded 2-dimensional modal logic from the more familiar 1-dimensional kind can be iterated: truth can be triply relativized to a *real* actual world $w - 1$, a ‘floating’ actual world w_2 , and a floating reference world w_3 .” (1980: 10)

³³One especially clear exposition of Frege’s ‘Sinn’-regress is Terence Parsons (1991).

³⁴Cf. (CM: 136 et seq.) Modal rationalism denies the existence of what Chalmers calls *strong* (metaphysical) necessities, i.e. truths the a posteriori necessity of which “can stem from factors

Modal empiricism: Take a borderline example, the supposedly necessary truth that cats are animals. We can certainly imagine a situation in which we would find out that some things we call “cats” are automata. This is *prima facie* evidence that we can imagine that cats are not animals.³⁵ But we cannot, given that cats are necessarily animals, imagine of cats that they might be animals. Chalmers’ reply would be that we are imagining is that catish things are automata. To imagine *this*, however, is much easier than to imagine (real) cats that are automata. The first task turns on what we can imagine about future technology, the second on what we believe to know about the essence of cats. What we *are* imagining depends on whether cats are essentially animals. So whether or not it is conceivable that cats (*real* cats, not just cat-like things) are automata depends on what biology tells us.

Is it, after all, really possible to conceive of a world where water is not H₂O? I am not so sure. No way of considering Twin Earth will make it into a world where water is not H₂O. And how can imagining a situation in which *p* is false count as a way of having it appear to one that *p* could have been *true*?

There are some important loose ends, however, which we would like to investigate in more detail in a longer paper: There is, first, the problem how primary intensions relate to what Chalmers has come to call “modal rationalism”. Against this, we would like to advocate something like ‘modal empiricism’, taking seriously the fact that not only what we mean by our words, but also what we find conceivable or even possible depends on what we know about the world we take to be actual. Another pack of problems contains the epistemology of possibility: could we consistently be, e.g., realists about the the inner, but fictionalists about the outer sphere of possibility? A third loose end concerns the (usually tacitly presupposed) connection between apriority and something like analyticity-in-the-language-of-thought.³⁶ Should we not, with Perry, apply “a priori” not to beliefs (or propositions or concepts), but to the *ways* in which we have them? Lots of open questions, then.

8 Conclusions

The dilemma I tried to set up for Chalmers shows that Two-Dimensionalism does not only keep what is right in Frege, but also inherits some of his problems.³⁷ Nothing can at the same time give the meaning and fix the reference, not even primary intensions can.

quite independent of the semantics of the terms involved”. Chalmers has another notion of ‘strong necessity’, according to which any a posteriori necessity is strong which has a necessary primary intension. As I do not think that primary intensions are an apt tool to explain aposterioricity in general, I do not use this notion.

³⁵Cf. Tidman: “It certainly *seems* we can imagine cats turning out to be automata. One could, say, imagine this about a particular cat in one’s lap. You might imagine finding a cleverly concealed zipper in the very fur you are stroking, which, when unzipped, reveals hidden riches of intricate computer circuitry. It seems wildly implausible to say that one is not thereby really imagining this of the cat.” (1994: 304-305)

³⁶Kit Fine has remarked that analyticity should be relativized to words: “Bachelors are unmarried”, e.g., is *analytic in* “bachelor” for it is true under all admissible definitions of that word. It seems that we need a similarly relativized notion of “a priori”. For whether or not a proposition will qualify as a priori will typically depend only on some of its component concepts.

³⁷It seems to us that Chalmers’ confidence was premature: “... a little looseness around the edges of a primary intension is entirely compatible with my applications of the [two-dimensional] framework.” (CM: 365)

9 References

- CDE Block, N. and Stalnaker, R. (1998), *Conceptual Analysis, Dualism and the Explanatory Gap*
<http://www.nyu.edu/gsas/dept/philo/faculty/block/papers/ExplanatoryGap.html>
- CH Byrne, A. (1998), *Cosmic Hermeneutics*,
<http://web.mit.edu/abyrne/www/CosmicHermeneutics.html>
- CAP Byrne, A. (2000), *Chalmers' Two-Dimensional Argument Against Physicalism*
- CM Chalmers, D. (1996), *The Conscious Mind*, Oxford University Press
- CC Chalmers, D. (1998), *The Components of Content*,
<http://www.u.arizona.edu/~chalmers/papers/content.ps>
- MMM Chalmers, D. (1998), *Materialism and the Metaphysics of Modality*,
<http://www.u.arizona.edu/~chalmers/papers/modality.html>
- MM Chalmers, D. (1998), *Mind and Modality*,
<http://www.u.arizona.edu/~chalmers/papers/mm.html>
- TS Chalmers, D. (2000), *The Tyranny of the Subjunctive*,
<http://www.u.arizona.edu/~chalmers/papers/tyranny.txt>
- CP Chalmers, D. (2001), *Does Conceivability Entail Possibility*,
<http://www.u.arizona.edu/~chalmers/papers/conceivability.html>
- SI Chalmers, D. (2001), *On Sense and Intension*,
<http://www.u.arizona.edu/~chalmers/papers/intension.html>
- CA Chalmers, D. and Jackson, F. (2000), *Conceptual Analysis and Reductive Explanation*,
<http://www.u.arizona.edu/~chalmers/papers/analysis.html>
- NE Chalmers, D. (2001), *The Nature of Epistemic Space*,
<http://www.u.arizona.edu/~chalmers/papers/espace.html>
- CC Yablo, St. (1998), *Concepts and Consciousness*,
<http://www.mit.edu/~yablo/concon.html>
- TK Yablo, St. (2000), *Textbook Kripkeanism & the Open Texture of Concepts*,
<http://www.mit.edu/~yablo/tk.html>
- CW Yablo, St. (2001), *Coulda, Woulda, Shoulda*,
<http://www.mit.edu/~yablo/coulda.pdf>
- F Fine, K. (1994), *Essence and Modality*, in: Tomberlin, J.E., ed., *Philosophical Perspectives 8: Logic and Language*, Ridgeview,
www.nyu.edu/gsas/dept/philo/faculty/fine/papers/essence.pdf
- P Prior, A.N. (1969), "The Possibly-true and the Possible", *Mind* 78, pp. 481-492, reprinted in: Prior, A.N., *Papers in Logics and Ethics*, Geach / Kenny, eds., University of Massachusetts Press, 1976, pp. 202-214
- AS Stalnaker, R. (1978), *Assertion, Syntax and Semantics* 9, reprinted in: Stalnaker, R. 1999, *Context and Content*, Oxford University Press
- IN Stalnaker, R. (1999), *Introduction*, in: Stalnaker, R. 1999, *Context and Content*, Oxford University Press
- CA Stalnaker, R. (2001), *On Considering a Possible World as Actual*,
www.nyu.edu/gsas/dept/philo/courses/content/papers/stalnaker.pdf
- T Tidman, P. (1994). *Conceivability as a Test for Possibility*, *American Philosophical Quarterly* 31, pp. 297-309