

Might water not be watery?

Philipp Keller*

June 30, 2001

Draft
comments welcome

Abstract

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-wordly 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: the "might" in "water might not be H_2O " is *merely* epistemic and represents "a [present] state of ignorance or uncertainty" (NN: 307, 103).¹ 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 ", [$\text{water is } H_2O$]₂, is the

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

¹It seems incorrect to me to say, as Chalmers does, that "Kripke allows that it might *turn out* that Hesperus is not Phosphorus" (SI: 17). He allows only that it *might have turned out* in a world in which we had not discovered that it could not.

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

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. 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”. According to Chalmers, what we should say not to misdescribe Twin Earth is that it is a world where watery stuff has a different microphysical structure than it has on our world.

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.

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 conditionals $p \rightarrow q$, where p is the conjunction of all physical truths and q is a particular phenomenal truth, 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.

The crucial step here is that any necessitation of $[Q]_1$ gives us an a priori entailment of 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. By this, we conclude from

- (4) There is a p -world considered as actual where q is false.

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)

that

(5) There is a p -world not in $[q]_1$.

The step from (4) to (5) is the crucial step in the 'conceivability implies possibility' argument: if we can conceive (positively, ideally, after long enough rational reflection etc.)⁵ of a world falsifying q , there is a sense in which q is contingent. There are only illusions, but no hallucinations of possibility according to Chalmers.

This crucial step relies on the conceptual equivalence of q and $[q]_1$. The argument for this conceptual equivalence is quite general and consists in a supposedly canonical method of explaining away illusions of contingency: Whenever someone falsely thinks that " $a = b$ " is contingent, he mistakes it for " $[a]_1 = [b]_1$ " 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.⁶ 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 imagining the meanings of our words to be different from what they actually are. Chalmers, however, wants "watery stuff" to do more than just to indicate that we might talk another language than we actually. 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.⁷ Chalmers, however, thinks that two sentences really are conceptually equivalent and differ only in that the first, but not the second, is metaphysically necessary. "Water" and "watery stuff" are conceptually equivalent because they are connected by the following bridge principle, which, supposedly, 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 phenomenal 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

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.

⁶"If any phenomenon is picked out in exactly the same way that we pick out pain, then that phenomenon *is* pain." (NN: 153, 340)

⁷Polywater, however, allegedly discovered in the Soviet Union, is described by Kripke as water which is not watery (NN 129, 323).

out the same thing, *a priori* identical, thereby basing his argument on a *analogy* with the water case.

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 three arguments for this, the first *ad hominem*, a internal second one and the third programatic. Here is the first: If primary intensions were not expressible, Chalmers 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.⁸ The internal argument concerns our motivation to introduce primary intensions in the first place, i.e. to explain away the illusion of contingency. Diagonalization is something we do in order to make sense of beliefs we take to be necessarily false. It should therefore give us something we can *ascribe* to people to which we do not want to ascribe belief in the impossible proposition. 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' 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”. For any world *w*, what water would be if *w* turned out to be actual is the watery stuff in *w*. We may ask whether, in order to establish coreferentiality in all possible worlds, 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 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 reference-fixing from a God's eye point of view, that 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.

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

(8) †(watery stuff) is watery stuff.

⁸If all the physicalist has to do is to come up with a function, his task is too easy. Chalmers' challenge is to the physicalist is to come up with some *X* such that water/watery stuff = pain/*X*.

where \ddagger is a rigidifying operator (Stalnaker’s ‘upside-down dagger’) that takes a description to form a singular term denoting in all possible worlds whatever uniquely satisfies the description in the actual world. “ $\ddagger(\text{the } \phi)$ ” is, in Kaplan’s terms, “a rigid description which induces a complex ‘representation’ of the referent into the content” (1989: 580)

I think one of the reasons why the contingent a priori has puzzled many philosophers is that claims that something is contingent and a priori are *doxastically indefensible*. If I claim to know (and even to know a priori) that p and at the same time admit that p might be false I do not contradict myself, but I will hardly convince anyone that I am right. I think that interesting a priori claims would have to be doxastically defensible. I will thus try to show that in both horns of the dilemma, we can evaluate (6) as a priori only if we assume that it is true. In the last part, I will investigate some consequences of this.

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”.

There is a chemical substance, heavy water or deuterium oxide, which is watery stuff, but not water, for it has the microstructure D_2O . It is odour- and tasteless and you could not tell it from water. So “water is watery stuff” is true only if “watery stuff” is taken to include a descriptive component excluding D_2O . As D_2O has been synthesized for the first time in 1933 and water has been discovered to be H_2O already in 1784, the description by which the reference of “water” actually 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. We can stipulate “watery stuff” to be *our* reference-fixing description. The claim, then, is that (8) is a case involving 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 the following is a priori but contingent:

(9) Julius invented the zip.

But there is a problem here, even granted that “the inventor of the zip” picks out a unique person in the actual world: We have no assurance whatsoever that this person invented the zip. 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. 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.

But we *stipulated* “Julius” to be a name of the inventor of the zip, whoever it was who invented the zip. The crucial question is: what does this stipulation amount to? 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. Our stipulation

now either 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 has to amount 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 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. By keeping the language fixed, then, we exclude worlds w where “Julius” does denote somebody other than whoever invented the zip in w . What is true of any world we may consider as actual is thus not (9) but the following conditional:

(10) If “Julius” denotes the inventor of the zip, then Julius invented the zip.

(10), 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 , namely the actual inventor of the zip, but in fact have named b .

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 H_2O that it is watery, that is water: you are supposing that “water” in the counterfactual world applies to whatever has the same macrophysical properties than H_2O in the actual world.

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 (??), “ $\phi(\ddagger(\text{the } \phi))$ ” (“whatever is ϕ is ϕ ”) is a truth of logic and can thus plausibly be taken to be a priori. Here are the matrices:

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.

“water”	H_2O	H_2O	H_2O	“watery stuff”	H_2O	XYZ	ABC
	XYZ	XYZ	XYZ		H_2O	XYZ	ABC
	ABC	ABC	ABC		H_2O	XYZ	ABC

Given these matrices, we indeed have $\text{water} = \ddagger(\text{water})$ and $\text{watery stuff} = \ddagger \text{water}$. But wait. How can we be sure that we have written down the right matrix for “watery stuff”. We stipulated that “watery stuff” should express the diagonal intension of “water” in the actual world. This was the projection of the diagonal onto the

horizontal. But then we did another step, namely we applied Davies' and Humberstone's "fixedly"-operator and replicated the upper row. Thereby we assumed that it does not matter for the evaluation of "watery stuff" in which world it is uttered. But perhaps it does. Let us assume that "watery stuff", if uttered in Twin Earth, 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 water concept (that we may express by "twatery stuff"). So we get the following matrix:

"water"	H ₂ O	H ₂ O	H ₂ O	"watery stuff"	H ₂ O	XYZ	ABC
	XYZ ₁	XYZ ₁	XYZ ₁		H ₂ O ₁	XYZ ₁	ABC ₁
	ABC ₂	ABC ₂	ABC ₂		H ₂ O ₂	XYZ ₂	ABC ₂

At first sight, this does not seem to change very much. Both concepts still have the same primary intension and so are a priori equivalent. If we assume that $XYZ = XYZ_1$ and $ABC = ABC_2$, it is still true that water = \dagger watery stuff. What we lost, however, is $\dagger \text{watery stuff} = \text{watery stuff}$. The new matrix for \dagger watery stuff = $\dagger\dagger$ water is the following:

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

Does this matter? I think it does. To see why, consider Davies's and Humberstone's rendering of the claim that (??) is a priori:

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

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 $\Box\dagger$, where \dagger is a function mapping the diagonal proposition of a propositional concept onto the horizontal. In a more uniform notion, this amounts to the following:

(12) $\Box\dagger(\text{water} = \dagger\text{water})$

(13) $\Box(\dagger\text{water} = \dagger\dagger\text{water})$

Though both "watery stuff" and " \dagger watery stuff" express the same diagonal proposition, the apriority claim only goes through if they do not differ. What does it mean to assume that they do not differ? It means that "watery stuff" is not only supposed to pick out the primary intension of "water" in every possible world *considered as counterfactual* but also in every possible world *considered as actual*. This is even stronger than the claim that their secondary intensions are the same. It amounts to the thesis that we cannot even conceive of a world *as actual* where the counterfactual behaviour of "watery stuff" and the primary intension of "water" differ, where e.g. $ABC_2 \neq ABC$. It has to be a priori that their secondary intensions are the same.

This, I think, is an unreasonable strong contention. Whatever our specification of the primary intension of "water", it seems that we can always sensibly ask whether we managed to capture it. Whether or not we managed to capture it, is an empirical question: it depends on what *we*, competent speakers of our language, would say in counterfactual circumstances. 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.

Here is another argument to the same effect: 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 .⁹ 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 by the inhabitants of w , where the substance they 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.

Whether or not “watery stuff” and (what we express as) “twatery stuff” have the same secondary intensions, i.e. pick out the same stuffs in all possible worlds considered as counterfactual, is an empirical issue. Part of this question concerns counterfactual language use: what competent speakers would say if such-and-such a world would turn out to be their actual world. Another part, however, concerns *actual* language use: whether my word “watery stuff” and what I express as “twatery stuff” apply to the same Twin-Earthian substance. “Twatery stuff”, by stipulation, applies to all and only to things having the same macrophysical properties than XYZ. But the only thing I know of the macrophysical properties of XYZ is that they are the same than those of H_2O . So to answer the question whether “watery stuff” and “twatery stuff” have the same secondary intensions, I have to rely on my knowledge that “watery stuff” describes all and only the macrophysical properties of H_2O , that is on my *de re* knowledge of water that it is watery.

This, however, means that “water is watery stuff” is a priori if it is true.¹⁰

6 Morals

The moral of both horns is, I think, that we have to keep fixed not only the language, but the whole actual world, in order to study the other-wordly behaviour of our words. 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. What our words mean, even in counterfactual circumstances, depends on how the world is.

But has Chalmers not argued for his claim that knowledge of primary intensions

⁹Imagine w to be future earth, where tradition is highly respected, but water has become sparse and is kept in laboratories, while another substance, say petrol, fills the lakes. It is not important whether you would describe this case as a case where the word “water” would have different properties, e.g. would be Millian or change from a descriptive to a proper name. What is important, however, is that it would still be the same word and nothing seems to go against this.

¹⁰In 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.

is a priori? Here is the argument:

“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.¹¹ 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.¹² “Conditional knowledge”, as Chalmers calls it, is not knowledge of a whole array of conditionals. It is true, of course, that sufficient descriptive knowledge about a world w allows you to identify the referents of your terms in w if everything does. But this does not mean that you know to what your terms would refer to if w turned out to be actual.¹³ To know this, you need further information, not about w , but about which of the worlds you consider as actual *is* actual. This is so even if w *is* your actual world. Even then you need, over and above the totality of truths in w , the further information that w is your actual world, in order to be able to take the meaning of “watery stuff” at face value. It is not just indexical information that you need, where and who you are, what time is it and the like. It is indexical information of the form “This world is my world” and the reliance on such information is just what makes a knowledge claim empirical.

The point can be illustrated by considering a disanalogy between the following two claims:

- (14) If p is true, p is necessarily true.
- (15) If p is true, p is a priori.

Kripke argues that we know the truth of (14) a priori, by philosophical analysis (NN 109, 311). But what about (15)? 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 .¹⁴ There is another difference: (14) makes necessity-with-respect-to- p automatically transitive. But apriority, even with respect to sentences like “water is watery stuff”, is not. It does not follow from the fact that you can conceive that you can conceive that p that you can conceive that p . What

¹¹Cf. Chalmers and Jackson: “Possession of a concept bestows a *conditional ability* to identify the concept’s extension, given information about hypothetical epistemic possibilities (in the broad sense of “epistemic possibility”, invoking hypotheses about the actual world that are not ruled out a priori).” (CA: 5)

¹²Note 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.

¹³Chalmers seems to think it does: “Like an epistemic intension but unlike a subjunctive intension, a two-dimensional intension can be evaluated a priori. One needs no empirical information about the actual world, since all the relevant information is specified in the epistemic possibility.” (SI: 21)

¹⁴“A *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)

you can conceive of, depends on your epistemic situation, your cognitive makeup and many other things (even what you can ideally conceive of does). You can conceive of these parameters being different, without, of course, thereby conceiving of what you can not conceive of.

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 this second option.

7 *n*-Dimensionalism

Even if every metaphysically possible world can be considered as actual, not all of them can be so considered at the same time and by the same person. You have to keep one world fixed and evaluate possible world candidates *for this world*. It might be that some of these candidates have more and different candidates than the world you keep fixed.¹⁵ But you can, as we did in the case of the primary intension of “watery stuff”, keep track of your commitment. The price to pay is a regress.

Recall how we got to the third dimension: we asked ourselves how we know that “watery stuff”, supposed to express the primary intension of “water”, does what it was supposed to be. Could we not, if we *were* inhabitants of Twin Earth, have another concept of water? So we indexed the secondary intension of “watery stuff” and, by diagonalization, got *its* primary intension. With respect to the expression we used to express this primary intension, namely “†(watery stuff)”, we simply assumed that its secondary intension is constant across worlds considered as actual. We thus implicitly applied the rigidifying operator \mathcal{F} . But in this we could have been wrong. To keep track of this possibility, we would need a fourth dimension. So we get some sort of a Fregean regress of higher-order senses.¹⁶

The Fregean regress is due to the presence of a functor $[\dots]_1$ mapping expressions to primary intensions in the object language and the claim that every primary intension should, at least in principle, be expressible.

Here is another way to argue for the induction step: A statement p is conceivable (i.e. $\neg p$ is not a priori), Chalmers claims, iff its 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, when we conceive of a p -world as actual, is not just conceiving of p having a non-empty primary intension. Instead, it 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:

$$(16) \quad p \text{ is conceivable} \iff “[p]_1” \text{ is possible}$$

¹⁵Note that the claim is that you could, if one of these worlds turned out to be actual, consider more and different worlds *as actual*. This is something Chalmers and Jackson miss in their reply to the argument that “If watery stuff is H₂ then H₂O is water” is not a priori because it would be false if uttered on Twin Earth: “The latter [“If H₂O would be watery stuff, then it would not be water” - I changed the example] is a subjunctive claim about an explicitly counterfactual scenario, so at best its truth reflects negatively on the *necessity* of an entailment from [“H₂O is watery stuff”] to [“Water is H₂O”].” (CA: 13)

¹⁶This has been remarked by Davies and Humberstone: “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)

By the definition of primary intensions, conceiving of “[p]₁” is conceiving of something that expresses *its* primary intension, “[[p]₁]₁”. So (16) gives us:

$$(17) \quad p \text{ is conceivable} \iff "[[p]_1]_1" \text{ is possible}$$

To conceive that p , it is not enough to believe that p expresses some truth or other at some world. One has to imagine that the sentence, in a world where its relevant meaning is preserved, expresses a truth which is conceptually equivalent to the truth we wanted to conceive of. 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. So one has to know its primary intension. In order to know that our evidence for the possibility of p does not just depend on our ignorance of its primary intension, 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. 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. And so we get the regress.

Here is an example: To answer the question whether “water is watery stuff” is a priori, we have to do the following: fix a world w we take to be the actual world, consider a world v as actual and ask ourselves whether what we would call “water” in v is what we call “watery stuff” in w , when both are evaluated in v . If what we *here* call “watery stuff” is not what we *there* would call “water”, then it is not a priori. In the preceding sentence, “here” and “there” function as operators on contexts of utterance, something Kaplan calls ‘monsters’. For we can vary the world we fix as actual and imagine the meaning of “watery stuff” to be different from what we take it to be.

Two days ago, we met an operator that could do that, namely “if I were you”. It seems to me that the fact that “If I were you then I would stay away from me” is sensible, but that “If I am you then I stay away from me” is not, is evidence that we can and do consider contexts of utterance as counterfactual. If this is true and if we accept Chalmers linguistic criterion for the ‘considering as actual’ vs. ‘considering as counterfactual’ distinction, then we can consider as counterfactual the possibility that the primary intensions of our words are different from what they in fact are.

We easily get other monsters. One is the operator studied by Davies and Humberstone ((1980: 10):

$$(18) \quad (w_1, w_2, w_3) \models R\alpha \iff (w_1, w_1, w_3) \models \alpha$$

which gives us strong necessities in the two-dimensional framework: $RA\alpha \leftrightarrow A\alpha$ will be deeply contingent (i.e. its \mathcal{FA} -modalization is false) but arguably a priori. But of course we can define another notion of a priori for the three-dimensional case.

Earlier in the talk we met another multidimensional operator, namely Stalnaker’s dagger, without adjunction of the “fixedly”-operator \mathcal{F} , corresponding roughly to the locution “what is said in the utterance s is true” and projecting the n -dimensional diagonal onto the $n - 1$ -dimensional hyperplane. It seems to me that iterations of this operator make sense.

I think that n -ary intensions could be useful in epistemology in general. Think of the world where you are a brain in a vat and of Putnam’s famous quick-and-dirty argument that if he were a brain in a vat, he could not consider the world as actual where he is a brain in a vat and so, necessarily, he is not a brain in a vat. This is hardly convincing, but why? Partly because he could reason *like this* (or at least utter the same words) if he *were* a brain in a vat. But even if he were a brain in a

vat, he could say”what would be said by “I am a brain in a vat” if I were not a brain in a vat is true”, and what he would say by this utterance would be true. Or take fictions: what is fun about fictions is not only that I can consider the fictional world as actual, but that I can imagine it *to be* actual, e.g. that I can take the fictional “I” to refer to me. For what I say when I specify what is said by your utterance of p does not only depend on what your utterance means but also what the words mean *I* use to specify the content of your utterance.

I hope to have laid out some middle ground between two extreme views on the a priori. On the one hand, we have Chalmers defending a substantive notion of apriority, keeping the language fixed and varying nevertheless in *some* sense at least the meaning of “water”. On the other hand, we have people who think that for p to have a necessary primary intension is for p to be true whatever it means. I think it is worth the effort to vary the meanings of words independently, to make distinctions between varying degrees of keeping the language fixed and thus to go along with more dimensions than just with 2.