

Christian Nitz 2004 Universität Bielefeld	published: <i>Christian Nitz: Two-Dimensionalism and Natural Kind Terms, Synthèse 138, 2004, 125–148</i>
--	--

Two-Dimensionalism and Natural Kind Terms

Christian Nitz
cnitz@uni-bielefeld.de

Kripke and Putnam have convinced most philosophers that we cannot do metaphysics of nature by analysing the senses of natural kind terms – simply because natural kind terms do not have senses. Neo-descriptivists, especially Frank Jackson and David Chalmers, believe that this view is mistaken. Merging classical descriptivism with a Kaplan-inspired two-dimensional framework, neo-descriptivists devise a semantics for natural kind terms that assigns natural kind terms so-called ‘primary intensions’. Since primary intensions are senses by other names, Jackson and Chalmers conclude that we can and should do metaphysics of nature by analysing the natural kind concepts competent speakers possess. I argue that neo-descriptivism does not provide a suitable basis for doing this kind of metaphysics. I first of all give a detailed account of the neo-descriptivist semantics and deflate the intuitive support neo-descriptivists try to draw from their case of the XYZ-world. I then present three arguments – the *Argument from Ignorance*, the *Argument from Conceptual Analysis*, and the *Argument from Laziness*. Taken together, these arguments undermine the neo-descriptivist analysis of natural kind terms. I conclude that natural kind terms do not have senses, that we cannot do metaphysics of nature by analysing the senses of our kind terms, and that the Kripke-Putnam account still provides the best semantics for natural kind terms we have.

1. Neo-Descriptivism and Armchair Metaphysics of Nature

Sometimes new trends in semantics come in, as it were, through the backdoor. The latest case in point is the account I will call ‘neo-descriptivism’. Even though neo-descriptivism is a stance in semantics, it has been developed within the metaphysical project propounded by David Chalmers and Frank Jackson (cf. Jackson 1998, esp. ch. 1-3; Chalmers 1996, esp. ch. 4; Chalmers & Jackson 2001). At the heart of this project lies the conviction that metaphysics is pursued in order to locate higher-order phenomena relative to the basic layer(s) of the world.¹ Chalmers and Jackson believe that it is the job of the metaphysician to reductively explain for instance how a fundamentally physical world can contain psychological facts, how a basically colourless world allows

for the application of colour-predicates, or how a fundamentally non-normative world can be the object of moral judgement.

The view that doing metaphysics comes down to locating higher-order phenomena is hardly peculiar. What is peculiar about the approach at hand is the idea that locating phenomena proceeds via *a priori* deductions. Jackson and Chalmers hold that any reductive explanation had better consist in a deduction of the higher-order facts from the fundamental account. They moreover think that any such deduction has to be *a priori*. Jackson and Chalmers of course acknowledge that the account of the fundamental layer will typically consist of empirically ascertained facts. But they believe that the inferences from the fundamental facts to the higher-order phenomena have to rest exclusively upon *a priori* available conceptual knowledge (cf. Chalmers & Jackson 2001). Hence, given that she knows the physical facts, all our metaphysician can rely upon to reductively explain psychological or moral phenomena is her grasp of our psychological and moral concepts.

Conceptual knowledge thus becomes the prime resource as well as the prime limiting factor for doing metaphysics. Chalmers and Jackson assure us that this is nothing to worry about. They are convinced that we can devise reductive explanations for moral and architectural as well as biological, astronomical, and chemical facts (cf. Chalmers 1996, 64, 73; Jackson 1998, ch. 4 and ch. 5; Chalmers & Jackson 2001, §§1-4). They for instance hold that we can *a priori* deduce ordinary truths about water from a scientific description of the world. Assuming that the scientific description contains propositions such as “H₂O covers most of the Earth” and “H₂O is the transparent, colourless etc. liquid of our acquaintance”, Chalmers and Jackson argue that we can infer that water covers most of the Earth. All we need to invoke is that water is the transparent, colourless etc. liquid of our acquaintance. And this, they hold, we know *a priori*.

To some the idea that metaphysics of nature can be done in the armchair might seem attractive. Still, the claim that we can *a priori* deduce facts about natural kinds from a fundamental physical description of the world is rather contentious, since it presupposes that we have substantial conceptual knowledge about natural kinds. This runs counter to established philosophical wisdom, for one might want to point out that Putnam and Kripke established that all substantial truths about natural kinds must be *a posteriori*. It is here where neo-descriptivism comes into play. The core idea of neo-descriptivism is that any semantic analysis must distinguish the conceptual dimension of an expression’s meaning from its metaphysical dimension. More precisely, neo-descriptivists think that the meaning of an expression can be factored into two intensions.² On the one hand, there is the expression’s primary intension.³ This intension captures the expression’s sense, i.e. its context-independent, purely conceptual content that is grasped *a priori* by the competent user and that quite generally determines the expression’s designation. On the other hand, there is the expression’s secondary intension. This intension captures the expression’s metaphysical content, i.e. the content the expression has given the actual context and given the expression’s sense. Since an ex-

pression's metaphysical content might be in part determined by non-linguistic aspects of the present context, it cannot in general be known *a priori*.

A neo-descriptivist semantics for natural kinds is precisely what Jackson and Chalmers need to support their idea that we can give *a priori* reductive explanations of ordinary truths about natural kinds. A neo-descriptivist account of natural kind terms implies that 'water', 'tiger', 'gold' and their ilk have primary intensions. Since these are senses, they are understood to yield substantial conceptual knowledge. Neo-descriptivism thus comes in handy to provide the very *a priori* insights that support the *a priori* deductions thought to be central to reductive explanations. It hence is hardly surprising that Chalmers and Jackson vigorously propound the neo-descriptivist enterprise. I am afraid that this is bad news. For I am going to argue that the prospects of analysing natural kind terms along neo-descriptivist lines are very dim indeed. In order to do so, I will proceed in two steps. Firstly, I will review in some detail the ideas of neo-descriptivism. In so doing, I will not worry too much about exegetical minutiae or terminological variations. I will rather present a somewhat idealised and unified account of neo-descriptivism (section 2).⁴ Secondly, I will present a number of arguments designed to show that neo-descriptivism shares the fate of classical descriptivism: it runs into severe trouble (section 3). I take this to vindicate the general line of the orthodox Kripke-Putnam account. I moreover take this quite generally to strengthen the case for a 'slim semantics', that is, for a semantics that strives to keep semantic and epistemological properties neatly apart.

2. The Neo-Descriptivist Picture

Neo-descriptivism is an eclectic enterprise. On the one hand, it takes up the main strands of the descriptivist tradition as well as the two-dimensionalist framework proffered by Kaplan and Stalnaker. On the other hand, it incorporates some of the main anti-descriptivist ideas put forth by Kripke and Putnam. Neo-descriptivism fuses these ideas into a new semantics by restricting the respective accounts to different semantic dimensions. Its core claims are these:

- (1) Classical descriptivism is right about natural kind terms. These terms do have descriptive senses, these senses are grasped by competent users, and they yield substantial *a priori* truths.
- (2) The Putnam-Kripke orthodoxy is right about the referential properties of natural kind terms. For instance, 'water' rigidly designates H₂O.
- (3) (1) and (2) can be squared once one presumes a generalised as well as modified two-dimensional framework.

Drawing on their distinction between primary and secondary intensions sketched above, neo-descriptivists propose a straightforward way to reconcile classical descriptivism with the Kripke-Putnam account. They claim that classical descriptivism is right about

the conceptual dimension of natural kind term meaning, whereas the Kripke-Putnam account is correct about the respective metaphysical dimension. According to neo-descriptivism, then, classical descriptivism and the Kripke-Putnam account aren't rivals at all.

The viability of this proposal depends on whether neo-descriptivists can come up with a plausible account of what primary and secondary intensions are. In order to explain this, neo-descriptivists embark on a manoeuvre similar to Kaplan's distinguishing two roles for possible worlds. Kaplan distinguishes contexts of utterance – i.e. the possible situations utterances are made in – from circumstances of evaluation – i.e. the possible situations utterances as made in contexts are evaluated at. Neo-descriptivists employ a similar distinction. They rely on the idea that we can take two different perspectives on possible worlds (cf. Chalmers 1996, 56ff; Jackson 1998, 46ff). On the one hand, we can conceive of possible worlds *as counterfactual*, regarding them as possible alternatives to our actual world or, to take up Kripke's phrase, as 'total ways the world might have been' (Kripke 1980, 18). This is the common way to think of possible worlds. In conceiving of a world as counterfactual, we hold on to our in fact actual world as our actual world. Hence, all semantic properties of all our expression will remain unaffected. This even holds of those terms whose semantic properties are determined externally, since they remain anchored in the very world they in fact are anchored in.

On the other hand, we can, neo-descriptivists maintain, conceive of a possible world *as actual*. That is, we can think of that world as being the actual world, pretending that it instead of our in fact actual world is the world we live in. If we think of a world that way, we have to acknowledge that all terms whose contents are determined externally might have different contents from those they in fact have. Even though one hence could take worlds considered as actual simply to be Kaplanian contexts, that is not the way neo-descriptivists think of them. Jackson and Chalmers rather understand worlds considered as actual to be 'epistemic possibilities', that is, 'specific way[s] the actual world might turn out to be, for all one can know a priori' (Chalmers & Jackson 2001, 324; cf. Chalmers 2002, §3.1). According to neo-descriptivists, this yields a subtle yet important difference: A world considered as actual thus understood is not bound to contain a token of an expression that is evaluated with respect to it (cf. Chalmers 2002, 20). But any context is.⁵

Drawing on these different perspectives on possible worlds, neo-descriptivists provide a characterisation of primary and secondary intensions. A primary intension is supposed to capture a sense, i.e. a reference determining and *a priori* knowable purely conceptual element in a term's meaning. Neo-descriptivists accordingly equate an expression's primary intension with a content that determines an extension for every possible world *considered as actual*. Primary intensions hence are modelled as functions $f: w_A \rightarrow e$ from worlds considered as actual to extensions. These intensions provide context-independent semantic specifications. Since worlds considered as actual are to be equated with epistemic possibilities, they moreover capture the 'epistemic dependence

of the extension of our expressions on the state of the world' (Chalmers 2002, 16). In order to grasp a primary intension, I have to know how the extension of the respective term depends on the way the world happens to be. That is something I can know *a priori*, since it does not presuppose any knowledge about my actual world.

If one conceives of primary intensions this way, one can very well maintain that the primary intension of a natural kind term is determined along the lines laid out by classical descriptivism. This is just what neo-descriptivists do. They for instance believe that the primary intension of 'water' is fixed by the role we assign the respective substance. Since we hold that water is a transparent liquid that quenches thirst, is odourless and colourless, falls from the sky as rain etc., the primary intension of 'water' is bound to contain these features. Neo-descriptivism thus yields an epistemic characterisation of the intensions of kind terms. Put succinctly, the primary intension of 'water' is the role we *a priori* know the substance in question to play. This role remains unaffected by variations of the respective actual world. Consider a world as actual that is just like ours, except for the fact that it contains the water-like substance XYZ where our world contains H₂O. Imagining this world to be our actual world does not change the role we *a priori* know water to play. All it does change is the stuff filling that role.

Secondary intensions, on the other hand, are designed to capture the metaphysical as well as possibly *a posteriori* dimension of a term's meaning. They are understood to be contents that determine extensions for worlds *considered as counterfactual*. Secondary intensions are hence to be modelled as functions $f: w_C \rightarrow e$ from worlds considered as counterfactual to extensions. A term's secondary intension provides a potentially context-dependent semantic specification: it spells out what the term designates, given its primary intension and given the actual world. In many cases, all external facts will be irrelevant. Terms like 'triangle' and 'bachelor' designate what they designate independent of matters contextual. A natural kind term's secondary intension, on the other hand, is externally determined by the stuff that satisfies the term's primary intension in the actual context, and it picks out *that* stuff in every world considered as counterfactual. In other words, the secondary intension of a kind term is the actual filler of the role as given by the term's primary intension. This explains why Kripke and Putnam are right. Since H₂O fills the water role around here, 'water' as uttered here rigidly denotes H₂O.

There are a few more important things to say about primary intensions. To begin with, it is important to mention that whenever we consider a world as actual, we view it as *centred*, i.e. with a location in that world that highlights a speaker, a time, a place, and a surrounding (cf. Chalmers 1995, 4ff; Chalmers 1996, 60f). For the extension of a term such as 'water' according to neo-descriptivist premises by the *local* environment rather than by the whole world – by the transparent liquid that is around here. Taking this into account, a term's primary intension must be represented formally as a function $f: w_A^* \rightarrow e$ from *centred* worlds considered as actual to extensions.

Secondly, neo-descriptivists hold that since the intensions of our terms might at first not be transparent to us, we have to employ the method of possible cases to get

clear about them (cf. Jackson 1998, 31-37; Chalmers & Jackson 2001, §3). According to this variety of conceptual analysis, what we are to do to determine the primary intension of, say, ‘water’ is this: we envisage a world and consider it as actual. We then draw on our intuitions to figure out what the term ‘water’ designates in that world. Doing this for a number of worlds will ideally yield a description ‘the stuff that satisfies F_1 & F_2 ... & F_n ’. This description spells out the role we *a priori* know water to play. It hence gives the primary intension of ‘water’. However, neo-descriptivists maintain that there is no guarantee that the process described will effect a finite description that precisely captures the expression’s primary intension. In many cases we have to make do with descriptions that at best approximate the term’s primary intension (cf. Chalmers & Jackson 2001, 322f). We still do know that the expression under scrutiny does have a primary intension, for we will find that speakers manifestly possess the ability to determine the term’s extensions in worlds considered as actual (cf. *ibid.*; Chalmers forthcoming b, 12, 28).

Finally, neo-descriptivists employ primary intensions to elucidate what it is for a sentence to be *a priori*. They maintain that a sentence s is *a priori* if and only if s is true in all worlds considered as actual. This so-called ‘core thesis’ (Chalmers 2002, 8) is first of all understood to state a condition of adequacy on any two-dimensional semantics that can satisfy the principles of a broadly Fregean semantics (cf. Chalmers 2002, §1; Chalmers forthcoming b, §5). We have to opt for an epistemic understanding of worlds considered as actual, neo-descriptivists argue, since what we want is a Fregean account and the contextual account doesn’t yield one, for it does not satisfy the core thesis (cf. Chalmers 2002). However, given that neo-descriptivism is committed to the core thesis, it yields a test for *a prioricity*: if s happens to be false in some world considered as actual, we apparently can safely conclude that s is not *a priori*.

The neo-descriptivist picture unfolded neatly fits with the metaphysical ideas sketched at the beginning. According to neo-descriptivism, we know *a priori*, say, that water is the colourless and odourless liquid that fills the lakes and falls from the sky etc.. Given suitable scientific results, we can hence *a priori* deduce from a scientific description of the world that water covers most of the Earth. This nicely locates our ordinary truths relative to the account of our world provided by the hard sciences. Hence, if neo-descriptivism provides a viable semantics for natural kind terms, the metaphysics of nature envisaged by Chalmers and Jackson appears to be feasible. Conversely, this metaphysics seems to be feasible only if the semantics proves to be correct, for this kind of semantics appears to be the only plausible source for the required substantial *a priori* truth about kinds. I have already indicated that I think that this is bad news for this kind of metaphysics. We are now going to see why this is so.

3. Assessing Neo-Descriptivism

I will argue that neo-descriptivism provides a skewed semantics for natural kind terms. In order to show this, I will present four arguments. The first is designed to deflate the intuitions that apparently support neo-descriptivism. The following three aim to show

that neo-descriptivism is flawed, concentrating on different aspects of the semantics proffered. Before I will come to this, some remarks are in order. Firstly, neo-descriptivists apparently believe that their semantics has merits even if it is considered in isolation from the metaphysical project that comes with it (cf. Chalmers 1995; Chalmers 2002). I will not discuss those. I will rather meet neo-descriptivism head on, arguing that it is the wrong theory for natural kind terms. Secondly, let me stress that there is nothing in general wrong with a two-dimensional semantics for terms other than indexicals. To the contrary, it is rather plausible that any general semantics will have to be two-dimensional, since it will have to deal with the pervasive phenomenon of context-dependence. The following arguments aim specifically at the epistemic two-dimensionalism proffered by Chalmers and Jackson. They do not affect a contextual understanding of two-dimensionalism that eschews the idea that there is substantial conceptual knowledge to be had; in fact, I think they support it.⁶ Finally, it would be wrong to claim that we have no *a priori* knowledge about kinds at all. One might very well want to hold that we do know *a priori* that, say, the stuff we picked out in introducing the term ‘water’ is necessarily water. Yet this knowledge is generated by mere stipulation and does not amount to substantial *a priori* knowledge about kinds.

3.1 Deflating Intuitions: The XYZ-World

The XYZ-world, as you will recall, is precisely like ours except for the fact that it contains a hypothetical water-like substance XYZ where our world contains H₂O. If we conceive of this world as actual, we pretend that it is *our* world, i.e., that we in fact are living in such a world (cf. Chalmers 2002, 17). Moreover, if we conceive of the XYZ-world as actual, it seems intuitively obvious that our term ‘water’ refers to XYZ. That is, if we set out to determine the extension of ‘water’ with respect to the XYZ-world, what we know about the distribution, behavior, and appearance of XYZ-molecules in that world will inevitably convince us that the extension of that term in the XYZ-world is XYZ. This might seem puzzling. For how can our expression ‘water’ possibly designate XYZ, given that, for all we know, ‘water’ rigidly designates H₂O?

Neo-descriptivists have an answer to this. They believe the sketched phenomenon to corroborate their claim that the designation of our term ‘water’ is determined by the term’s primary intension. By having grasped this intension, they maintain, we know that ‘water’ picks out the stuff that plays the role we *a priori* know water to play, which is why we intuitively and unanimously take ‘water’ to apply to XYZ, given that we conceive of the XYZ-world as actual. Neo-descriptivists thus maintain that the diagnosed change in designation is effected in a systematic way by the primary intension of our expression ‘water’. With respect to our in fact actual world, ‘water’, having the primary intension it does, rigidly designates H₂O, since H₂O plays the water role around here. With respect to the XYZ-world considered as actual, ‘water’, having the primary intension it does, rigidly designates XYZ, since this is what plays the water role in it.

However, there is a simpler explanation to be had: In the XYZ-world, my term ‘water’ means something different from what it actually means.⁷ This explanation

neatly resolves the problem at hand. It moreover fits nicely with the externalist assumption that the meaning of our term ‘water’ is in part determined by our environment. In fact, on externalist premises, it should come as no surprise that my word ‘water’ designates XYZ under the assumption that the XYZ-world is my actual world. We assume that the XYZ-world is my actual world, i.e. the world I do in fact live in. It hence is the world I acquire my term ‘water’ in. But if I acquire my word ‘water’ in the XYZ-world, I use a sample of XYZ to introduce that term rather than a sample of H₂O. (I hit on such a sample only if I employ the same introductory procedure for my term ‘water’ that I in fact employ; but I take it that this is part of the story.) This explanation allows for a simpler semantics. Since it diagnoses a humdrum change *of meaning* rather than a systematic change of designation *effected by an invariant meaning*, it can make do without building the resources for systematic changes in designation into the meanings of our natural kind terms. That is to say, it makes do without primary intensions.⁸

This explanation even squares better with our intuitions than the neo-descriptivist account. Intuitively, we are happy to diagnose changes in meaning to account for divergent referential properties. Just consider a world that is almost precisely like ours. Yet in this world, the just born Ernest Hemingway is taken from his cradle, murdered, and replaced with the son of a local beggar before someone in Oak Park, Illinois, had even thought of calling the poor child ‘Ernest’. This name is rather used for the impostor who becomes a world-famous writer, adventurer, and big game hunter. It is intuitively evident that if this world is our actual world, the name ‘Ernest Hemingway’ we use does not refer to Ernest Hemingway but rather to the gifted fraud. Yet in analysing this case, we do not even consider to inflate the semantics of ‘Ernest Hemingway’ along neo-descriptivist lines. We simply acknowledge that the name has a different meaning.

Here is another argument to the same effect. As everybody agrees, the semantic properties of, say, ‘elm tree’ are exhaustively determined by the total use we make of that term and the environment this use takes place in.⁹ Roughly, our term ‘elm tree’ designates elm trees because we call certain trees by that name and these trees happen to be elm trees. Moreover, it is hardly controversial that a different use can make for a different meaning. A community that is exactly like ours, save for the fact that its members call beech trees ‘elm trees’, would bestow a different meaning on the term. But if that is so, it is hard to see why differences in the environment cannot bring about changes in meaning as well. Imagine a world which is just like ours save that wherever there is a beech tree in our world there is an elm tree in it (and vice versa). If this is my actual world, I am calling beech trees ‘elm tree’. I hence do not employ the term with the meaning it has in the in fact actual world. Very much the same can be said about the XYZ-world.

The case of the XYZ-world, then, does not corroborate neo-descriptivism’s semantics, since it can very well be accounted for by an orthodox Kripke-Putnam account. That is to say, one can very well account for the fact that our term ‘water’ designates XYZ given that the XYZ-world is actual without invoking the assumption that the meaning of ‘water’ comprises an element — its primary intension — that determines

the term's designation. All one has to invoke is the externalist idea that the meaning of a kind term depends on the way it is introduced combined with the nature of the sample employed in that process. This hardly amounts to a knock-down argument against neo-descriptivism. Yet it does deflate its intuitive basis.

3.2 The Argument from Ignorance¹⁰

Neo-descriptivism builds a decidedly immodest epistemology into its semantics for natural kind terms. Neo-descriptivists maintain that anyone who understands natural kind terms must have grasped their primary intensions. Since these are determined by the roles we assign to kinds, it follows that any competent speaker is bound to have grasped the essentials about natural kinds. At first sight, this appears to be nothing to worry about. Most of us will not know, say, that gold belongs to the transition metals and melts at 1064.43° Centigrade. Yet most of us will know that gold is a malleable and very valuable metal, that it is in most cases yellowish, that it is mined in South Africa, is stored in Fort Knox, and was back in the 80s falsely believed to be a good investment. What is more, we evidently expect the members of our community to possess such knowledge. This is precisely what neo-descriptivism predicts.

However, a closer look reveals that the neo-descriptivist epistemology runs into severe trouble. To begin with, we do not need primary intensions to account for our expectations. We quite naturally presuppose that the members of our community know certain facts about the world. For instance, we quite naturally presuppose that they know that London is the capital of England and that the earth revolves around the sun. Hence, it is not at all puzzling that we expect them to know that gold is a mostly yellowish metal, even if they do not have to know this in order to be competent speakers. This explanation squares better with our actual behaviour than the one provided by neo-descriptivism. Just think of the way we treat the local ignoramus. We do not assume that his mastery of the English language is impaired because he neither knows that gold is malleable nor that gold is a metal. If he happens to be a usually well-informed stockbroker and tells us that the gold price is just about to plummet, we of course rush to sell our South Africa government bonds.

Secondly, neo-descriptivism's contention that understanding a kind term presupposes grasping the term's primary intension conflicts with an observation emphasised by the proponents of externalist accounts: ordinary speakers simply do not know much about natural kinds. This might go unnoticed as long as we are concerned with water, gold, or tigers. But if we think about, say, magnesium or tapirs, this becomes obvious. To be sure, almost everybody will know something about the metal and the animal in question, e.g. that magnesium is used in flares or that tapirs are four-legged animals with trunks. Yet almost nobody will be able to come up with an account that is rich enough to determine credible primary intensions for the terms. As neo-descriptivists maintain, such an account has to be purely qualitative (cf. Chalmers 2002, §3.5). It hence is not allowed to contain pieces of non-qualitative identifying knowledge some of us might possess, e.g. "Tapirs are the animals I've seen at Berlin Zoo that one winter

day”. What is more, such an account would have to be rich enough to single out all and only the tapirs in all possible worlds considered as actual. But it is unlikely that any purely qualitative account an ordinary speaker can come up with would even single out the tapirs in our world.

Finally, neo-descriptivism’s epistemology evidently cannot deal with speakers who endorse eccentric theories about certain natural kinds. Imagine John to believe that gold is actually a radioactive mineral from outer space, a fact most people are ignorant of since our governments are desperate to cover it up. If John now asserts “The US keep their gold in Fort Knox for good reasons”, neo-descriptivists have to deny that he just claimed that the US keep their gold in Fort Knox for good reasons. For the primary intension determined by his understanding and the primary intension determined by our account are very different indeed. They for sure do not pick out the same stuff in our world – on neo-descriptivist premises, our term ‘gold’ might very well pick out some of the gold that happens to be around, whereas John’s term has an empty extension. However, it seems to be fairly obvious that we can and of course would disagree with John’s bizarre theory. But that presupposes that our term ‘gold’ has at least roughly the same reference as John’s term ‘gold’. On neo-descriptivist premises, this cannot be the case.

To be sure, neo-descriptivists believe that these arguments do not affect their stance. For they acknowledge that different speakers might very well assign *different* roles to, and hence associate *different* primary intensions with, one and the same natural kind term (cf. Chalmers & Jackson 2001, 327; Chalmers 2002, 32; Chalmers forthcoming b, 30). They even hold that these intensions might be *very* different indeed: You, being a city-dweller who knows nothing of oceans, might use ‘water’ non-deferentially for the liquid that comes out of faucets, whereas I, being a beach-dweller who knows nothing of faucets, might use ‘water’ non-deferentially for the liquid in the oceans (cf. Chalmers & Jackson 2001, 328). The variability of primary intensions does not, they argue, undercut successful communication, and it does not forestall disagreement, since both can be grounded in the common referent (cf. Chalmers forthcoming b, 32).

This is hardly a convincing response, though. On the one hand, it does nothing to solve the problem of eccentric primary intensions. On neo-descriptivist premises, a term’s primary intension is what determines the term’s referent: ‘gold’ applies to whatever satisfies the role associated with it. But if that is so, there is no common referent that could ground disagreement between John and us, since there just is nothing that satisfies the eccentric role he propounds. On the other hand, acknowledging variability trades a serious problem for a very serious one. For, as neo-descriptivists agree, if primary intensions are subject-relative, so is *a priori* knowledge (cf. Chalmers & Jackson 2001, 327; Chalmers forthcoming b, 30) — some things might be *a priori* to me, but not to you. But if that is so, what is it for a sentence *s* to be *a priori simpliciter* and hence suited to license the *a priori* deductions needed for reductive explanations?

Firstly, one could hold that a sentence *s* is *a priori simpliciter* if it is *a priori* to some speaker in our community.¹¹ That, however, will yield far too many *a priori* truths. For if *P* describes a procedure to successfully identify an instance of a natural

kind n , there might be a speaker in our community for whom it is *a priori* that n satisfies P . For instance, there might be speakers who can, without recourse to experience, justify that alligators are dangerous or that water flows from faucets, since this is just how they non-deferentially use ‘alligator’ and ‘water’, respectively. Secondly, one could hold that a sentence s is *a priori simpliciter* if it is *a priori* ‘for any given subject and time in our community’ (Chalmers & Jackson 2001, 320)¹². This evidently yields too few *a priori* truths. For if primary intensions are allowed to vary as outlined, it is almost certain that for any purportedly *a priori* sentence s , there will be a speaker in our community who assigns to s a primary intension that does not hold in all worlds considered as actual. Thirdly, one could maintain that s is *a priori simpliciter* if and only if it is *a priori* for a speaker ‘given ideal rational reflection’ (Chalmers forthcoming b, 30). But it is hard to see how improved rational powers can change anything. For instance, assume that I am the beach-dweller mentioned above. For me, the only *a priori* truth about water is that water is the liquid in the oceans. I cannot see how ideal powers of rational reflection could possibly make me change my mind on this, leading me to richer — and intuitively more accurate — conceptual truths about water.

The neo-descriptivist manoeuvre hence does not solve the problem of ignorance. To the contrary, I suspect that it rather adds to the appeal of the rival Kripke-Putnam account. Allowing variability leads neo-descriptivists to admit that communication concerning kinds is not grounded in what we believe about them, but rather in the shared referent — just as Kripke and Putnam maintain. Moreover, given the difficulties arising from variability, one might very well be tempted to adopt the simple solution propounded by Kripke and Putnam who hold that beliefs such as ‘gold is that mostly yellowish, malleable, and valuable metal’ do not enter into the meanings of kind terms *at all*. All they do is pick out the samples that determine these meanings.

3.3 The Argument from Conceptual Analysis

Neo-descriptivists will be quick to diagnose a serious flaw in the argument from ignorance. Chalmers and Jackson stress that in order to grasp a primary intension, you do not have to acquire propositional knowledge (cf. Chalmers & Jackson 2001, 320-323). Therefore, worrying about ignorance — i.e., lack of propositional knowledge — appears to be beside the point. According to neo-descriptivists, a speaker’s grasp of a primary intension is rather exhibited in her manifested ability to identify extensions in worlds considered as actual. More precisely, neo-descriptivists hold that a speaker’s grasp of the primary intension of a term t consists in her having acquired ‘tacit criteria for identifying the extension of the expression’ (Chalmers forthcoming b, 7). These criteria can be gauged from the features the speaker’s intuitive yet reflective application of t in worlds considered as actual is sensitive to (cf. Chalmers & Jackson 2001, 322). That a speaker assigns primary intensions to her terms, and which intensions she does assign, is hence revealed by her ‘rational evaluation of specific epistemic possibilities’ (Chalmers forthcoming b, 12); she does not have to come up with any explicit definition, theory or description. In accordance with this, neo-descriptivists emphasize that

descriptions like ‘the dominant clear, drinkable liquid in the oceans and lakes’ (Chalmers 1996, 57) at best roughly approximate the primary intension of ‘water’.

Neo-descriptivists now believe that a glance at our intuitions proves that we can identify the extensions of natural kind terms in worlds considered as actual. It thus is obvious, they think, that we assign primary intensions to our kind terms; a good look at our intuitions will moreover allow us to see what they are. In short, neo-descriptivists hold that our intuitions reveal primary intensions. I think there are good reasons to reject this claim. I grant that we do have intuitions about how to apply ‘water’ or ‘gold’ in some worlds considered as actual. But I do not believe these intuitions to show that we assign primary intensions to our kind terms. On scrutiny, it rather turns out that our intuitions concerning the application of kind terms are not robust enough to license the ascription of primary intensions — or so I will argue.¹³ I will moreover argue that we can account for our intuitions without assuming that kind terms have primary intensions.

Consider a world that is almost exactly like ours. Just like ours, it contains a pure liquid that quenches thirst, falls from the sky as rain, can be used to extinguish fire, fills the lakes and oceans, comes in a glass with any espresso you order in any decent café etc.. However, assume that the liquid in question is greenish and non-transparent. If we conceive of this world as actual, we can no doubt identify the extension of ‘water’: ‘water’ in this world applies to the liquid mentioned. From this it follows that the feature *being transparent* does not guide our application of ‘water’ across worlds considered as actual: in some worlds, our application of ‘water’ is guided by this feature, yet in other worlds, it is not. Similar cases can be devised for (almost) any feature intuitively relevant to the application of ‘water’. Just imagine a pure liquid that falls from the sky as rain, fills the oceans etc., but does not quench thirst. Or imagine a pure liquid that quenches thirst, falls from the sky as rain, can be used to extinguish fire etc., but does not fill the local lakes and oceans. A scrutiny of our intuitions thus shows that for almost any feature *F* that guides or application of ‘water’ in some world considered as actual, there is some such world in which the stuff we intuitively single out as ‘water’ is not *F*.¹⁴

Our intuitions concerning the application of kind terms hence are not as smooth as one might think. It rather appears that our tacit criteria for applying a kind term *k* do, as it were, come as a disjunction. Roughly, we apply *k* to some stuff *x* in a world considered as actual *w* if and only if *x* has the set of features *A*, or *x* has the set of features *B*, or *x* has the set of features *C* etc., where the sets involved do not have to overlap that much.¹⁵ Anyone with neo-descriptivist leanings should be worried by this. For if you still want to hold that our intuitions reveal primary intensions, you have to bite quite a number of bullets. First of all, you have to admit that our intuitions concerning ‘water’ or ‘tiger’ are very different from our intuitions concerning, say, ‘knowledge’. In the latter case, we have no reason to believe that our tacit criteria for applying the term are disjunctive. We can even point out some features that reliably guide our application of ‘knowledge’ across worlds considered as actual: we are sure that any case of knowledge

must be a case of justified true belief. Our concept of knowledge thus has a conceptual core. The same is true of most other concepts commonly believed to yield conceptual truths. Hence, contrary to what neo-descriptivists maintain, (cf. Chalmers & Jackson 2001, 320-328), the fact that the method of possible cases might work for ‘knowledge’ does nothing to show that it will work for natural kind terms as well. One might even wonder whether there can be a concept at all if there is no conceptual core.

Secondly, you have to admit that there are hardly any simple *a priori* truths about natural kinds. According to the core thesis, ‘Water is *F*’ is an *a priori* truth only if it is true in all worlds considered as actual. Given that our criteria come as a disjunction, this is true only if *F* is included in every set of features *A*, *B*, *C* etc.. But as we have seen, almost no feature is.¹⁶ Consequently, almost none of those purportedly obvious *a priori* truths so dear to proponents of classical descriptivism turns out to be in fact an *a priori* truth. For instance, neither ‘Water quenches thirst’ nor ‘Gold is malleable’ is in fact *a priori*. Anyone drawn to neo-descriptivism hoping to secure seemingly humdrum conceptual truths must thus be disappointed.

Finally, you have to admit that the semantic values a neo-descriptivist semantics assigns to utterances of kind terms are absurdly cumbersome indeed. If the primary intensions we assign to kind terms are revealed by our tacit criteria, and if these criteria come as disjunctions, it becomes apparent that ‘the dominant clear, drinkable liquid in the oceans and lakes’ is no more a credible approximation to the primary intension of ‘water’ than ‘Pluto’ is a credible approximation to the extension of ‘planet of this solar system’. Consequently, the primary intension I express uttering “Plants need water” is not even roughly that plants are in need of the dominant clear, drinkable liquid in the oceans and lakes. It is rather bound to be something awkwardly disjunctive and hence absurdly more complicated.

Let’s have a look at another case. Imagine a world as actual that contains a pure liquid that is tasteless, transparent, fills the lakes and oceans, can be used to extinguish fire etc.. Now imagine that this liquid never freezes, never evaporates, and that it is such that we can walk on its surface. Is this stuff water? All we can invoke to decide this matter are our intuitions. But these, I reckon, simply do not support a verdict one way or the other. To be sure, we have the strong feeling that there is something wrong with a watery stuff one can walk on. This is to be expected. For our intuitions concerning kinds are at least in part informed by the law-like generalisations we believe to hold around here – and the latter tell us that you cannot walk on water. However, what we want is to tease out our concepts. But it will often be very hard to draw a clear line between conceptual aspects and nomological presuppositions, especially if the world we consider as actual differs nomologically from our in fact actual world. In any such case we are likely to end up without a verdict about the application of our kind terms.

Our intuitions concerning the application of kind terms hence are not as steady as one might think. If you nonetheless want to hold on to the idea that they reveal primary intensions, you have to admit that these intensions are indeterminate with respect to many worlds considered as actual. To be sure, neo-descriptivists grant this:

Determinate application (...) may be restricted to epistemic possibilities that are not too far from home. When epistemic possibilities deviate greatly from our presuppositions about the actual world, some of our concepts will lose determinate application. (Chalmers & Jackson 2001, 324, Fn 9)

Indeterminacy *per se* should not trouble you. Any believable semantics has to account for a reasonable degree of indeterminacy, since most of our concepts lack determinate application in some circumstances. What should trouble you, however, is the unusually strong as well as systematic indeterminacy neo-descriptivism gives rise to.

Please note for a start that the walking-on-water-world proliferates — it appears easy to devise a multiplicity of cases along its lines. It is, however, not that easy to devise scenarios where our intuitions concerning the application of, say, ‘red’, ‘triangle’, or ‘knowledge’ break down.¹⁷ The extent of indeterminacy that besets the purported primary intensions of kind terms hence appears to be unusually great. Very much the same is true of its range. We typically know whether or not to apply some non-kind-concept in some world considered as actual if the world in question is, intuitively speaking, well-behaved. However, the walking-on-water-world and its cognates are well-behaved: the items they contain are relatively stable, their laws allow for reliable predictions, and their histories comprise recurring patterns of events. The indeterminacy of kind terms hence affects even worlds that usually do not give rise to conceptual uncertainties.

Moreover, it appears that our ability to identify extensions for kind terms is limited in a systematic fashion: If *k* is a natural kind term, we very often cannot apply *k* beyond worlds considered as actual whose *k*-laws are very similar to the ones that hold in our world.¹⁸ But if that is so, one might first of all wonder how plausible it is to claim that we associate a concept with *k*, even though our application of *k* across that dimension of worlds especially devised to model our conceptual resources is severely restricted. What is more, neo-descriptivism seems at least bound to explain why the intuitive application of our kind terms is restricted the way it is. There is a straightforward way to do so: neo-descriptivists could maintain that the meanings of our natural kind terms are at least in part determined by our folk-theoretic laws, and that it is hence hardly surprising that they cannot be applied in worlds considered as actual where these laws do not hold.¹⁹ However, this would make your grasp of a natural kind concept once again dependent upon the theories you are acquainted with. Hence, this would make neo-descriptivism once again vulnerable to the argument from ignorance.

In spite of this all, you could still hold on to the idea that our intuitions concerning the application of kind terms reveal primary intensions. But you do not have to. There is a simpler explanation as to why our intuitions are as restricted as they are — or indeed, why we do have intuitions concerning the application of our kind terms in the first place.²⁰ Our behavior shows that we possess more or less tacit criteria for the identification of, say, water, gold, or tigers: we cautiously take sips from wells, carefully bite coins, or look out for black stripes on yellowish fur. These procedures can be taken to be ‘local epistemic shortcuts’, that is, simple reliable procedures to identify specimens

of kinds *in our in fact actual world*.²¹ Such procedures do certainly come in handy in dealing with the world, just like being able to tell a chardonnay from a pinot grigio does. Yet these procedures are not conceptual. They are *a posteriori*, if anything is: you do not need to acquire any such procedure in order to be a competent user of a natural kind term, and the shortcuts you know do not affect the meanings of your utterances involving kind terms. Still, our local epistemic shortcuts can account for the intuitions we associate with kind terms. Since they are designed to work in this world, and since their reliability depends on the law-like generalizations that hold around here, they will give rise to intuitions as to whether something is, say, water or gold in many worlds nomologically similar to ours. But they will not work beyond these.

Neo-descriptivists maintain that our intuitions concerning the applications of natural kind terms in worlds considered as actual reveal primary intensions. I have given reasons to be rather sceptical about this. If we look at possible cases, we find, on the one hand, that our tacit criteria for the application of our kind terms come as disjunctions. If you nonetheless endorse neo-descriptivism, you have to admit that natural kind terms do not have conceptual cores, that they do have absurdly complex semantic values, and that almost none of the purportedly simple *a priori* truths so dear to classical descriptivism is in fact *a priori*. On the other hand, we find that our intuitions concerning the application of kind terms do not yield a verdict for very many worlds considered as actual. If you nonetheless hold up neo-descriptivism, you have to admit that the primary intensions of our kind terms are subject to an unduly strong as well as systematic indeterminacy. Taken together, I take it that all this justifies a rather different conclusion. I would hence like to conclude that neo-descriptivism is wrong: our intuitions concerning the application of kind terms do *not* reveal primary intensions; they simply are not robust enough. This is not to deny that we have intuitions as to how to apply our kind terms in some worlds considered as actual. But as I have argued, these intuitions can be accounted for without assuming that we assign primary intensions to natural kind terms. It might come as no surprise that I consider the fact that the respective explanation is in line with the project of a slim semantics to be one of its virtues.

3.4 The Argument from Laziness

Natural kind terms, neo-descriptivists assure us, have primary intensions. These intensions do not only yield extensions for worlds considered as actual; they also determine the secondary intensions of our kind terms. According to neo-descriptivism, the secondary intension of a kind term *k* in some world considered as actual *w* picks what *k*'s primary intension singles out in *w* in every world considered as counterfactual. Hence, if our in fact actual world is the actual world, the secondary intension of 'water' picks out H₂O in every counterfactual world, since H₂O is what the primary intension of 'water' singles out around here. Yet if the XYZ-world happens to be the actual world, then the secondary intension of 'water' picks out XYZ in every world considered as counterfactual. Any neo-descriptivist semantics for natural kind terms thus builds the resources for systematic changes in secondary intensions into the meanings of natural kind terms. In

that respect, a neo-descriptivist semantics for ‘water’, ‘tiger’, and the like will be fairly similar to any ordinary semantics for occasion terms like ‘I’, ‘here’, or ‘yesterday’.

However, please recall what we do to conceive of a world as actual: we take one ‘specific way the actual world might turn out to be, for all one can know a priori’ (Chalmers & Jackson 2001, 324) and imagine that it captures the way our world actually is. That is, we imagine that this world is *the* actual world, the world we live in. But if that is what we do, it is plausible to assume that we cannot leave the in fact actual world — we are, as it were, stuck in this actuality. For there is just one specific way our in fact actual world is, even though we might not know all its details (and we for sure cannot know them by *a priori* reflection). In other words, I can imagine that my actual world is a different one, but my actual world will never be a different one. I guess neo-descriptivists are happy to concede this.²² So let me assume that we are indeed stuck in this actuality. But if that is so, neo-descriptivists have to admit that my term ‘water’ will always designate H₂O. For I will simply never get into any circumstances where the primary intension assigned to the term will make it change its designation; indeed, I cannot. Neo-descriptivists hence have to hold that even though our natural kind terms could in principle have different secondary intensions, they in fact cannot.

There hence is an interesting difference between a neo-descriptivist account of natural kind terms and any run of the mill semantics for occasion terms. Let us call an expression ‘lazy’ if its semantics comprises the resources for it to systematically change its designation or intension, but within all those circumstances we can possibly get into, it simply never does. On any ordinary semantics, occasion terms like ‘here’ or ‘tomorrow’ are far from being lazy; we can and do witness them to change their reference. On neo-descriptivist premises, however, natural kind terms turn out to be lazy terms. I think that this yields another reason to reject the neo-descriptivist account of kind terms. For it is rather implausible that our natural kind terms are indeed lazy expressions.

On the one hand, there is a general reason to doubt that natural languages will contain many lazy expressions. Natural languages might be rather disorderly affairs. But it is a sound rule of thumb to assume that natural languages are by and large economical, or at least economical within their respective general structure. That is to say, if an expressions has some semantic property *F*, it usually does so because its having *F* plays a role in communicating or in describing the world. This is in part due to the fact that natural languages evolve and that excrescent features simply vanish in due course. Hence, it is simply implausible to maintain that natural languages will contain more than a few lazy expressions. Since our language does contain an abundance of natural kind terms, it is hard to believe that natural kind terms should fall into that semantic category.²³ Moreover, it seems to be a matter of good methodology not to assign terms semantic properties that, as Evans has nicely put it, ‘allow them to get up to tricks they never in fact get up to’ (Evans 1979, 190). In other words, we should refrain from assigning terms semantic properties that are idle. But this is precisely what neo-descriptivism does.

On the other hand, natural languages are learned. Crudely though plausibly, learning a language includes getting to know the meanings of its expressions. If we stick to a rough outline of the complex processes involved, it is fairly easy to imagine how Kurt might come to know the meaning of, say, ‘George W. Bush’: He simply witnesses utterances of that name, and after a while he figures out that any utterance of ‘George W. Bush’ names the same individual, *viz.* George W. Bush. It also is fairly easy to imagine how Kurt might acquire the meaning of ‘I’. Witnessing utterances of that indexical, he will after a while figure out how the reference of ‘I’ varies with the context. Now imagine that Kurt never encounters an utterance of ‘I’ that does not refer to George W. Bush. Imagine that the only person interfering with Kurt’s learning history who ever uses ‘I’ happens to be George W. Bush. This gives Kurt no reason to believe that the reference of ‘I’ varies at all. He will rather think that ‘I’ is just another name for George W. Bush. Generalising, it appears plausible that someone will learn that a term’s designation might vary if she has reason to believe that she has encountered utterances of that expression with divergent referents during her learning history. But this can never be true of lazy terms, because their reference does not vary at all within our reach. We simply never encounter a situation in which ‘water’ does not refer to H₂O. We consequently have no reason to build the resources for a systematic change in designation into the meanings of this term. Hence, even if ‘water’ had a primary intension, one could not and would not learn that it did.

4. Upshot: Armchair Metaphysics and the Orthodoxy

We can, Jackson and Chalmers assure us, do metaphysics of nature in the armchair. For metaphysics is limited by conceptual knowledge only, and we do possess substantial conceptual knowledge about natural kinds. To sustain this view, they present a neo-descriptivist semantics for natural kind terms that combines the core ideas of classical descriptivism with the Kripke-Putnam idea that natural kind terms designate rigidly and refer directly. But as a semantics for natural kind terms, neo-descriptivism runs into severe trouble. It does so for the reasons similar to those that sealed the fate of classical descriptivism: It relies upon mistaken assumptions about what competent speakers know, and it assigns kind terms semantic properties that are neither substantiated by, nor square well with, our intuitions. What is more, an orthodox Kripke-Putnam semantics can by and large account for the intuitions we do in fact have in a more compelling manner. I thus conclude that neo-descriptivism, as it has been outlined here, does not provide an adequate semantics for natural kind terms. This neither shows that neo-descriptivism — or, for that matter, two-dimensionalism in general — is a flawed approach to semantics *simpliciter*; nor does it show any two-dimensional semantics for natural kinds to be mistaken. But it does demonstrate that any such semantics had better respect the core insight of the orthodox Kripke-Putnam account, *viz.* that natural kind terms do not have senses that give rise to substantial *a priori* knowledge. What is more, if conceptual knowledge is indeed the prime resource as well as the prime limiting factor for doing metaphysics Jackson-Chalmers style, it follows that we cannot do meta-

physics of nature that way. We rather have to acknowledge that reductive explanations of accounts concerning nature do rest upon scientifically ascertained *a posteriori* knowledge – just as Kripke and Putnam said they would.²⁴

Notes

¹ Cf. Chalmers & Jackson 2001; Jackson 1998, ch. 1. See also Jackson 1994, Jackson 1994b and especially Lewis 1994. Cf. also Chalmers 1996, ch. 2.

² This is the core thesis of two-dimensionalism. Varieties of two-dimensionalism are endorsed in Stalnaker 1978, Evans 1979, Lewis 1981, Chalmers 1996, ch. 2.2, and Jackson 1998, ch. 2. For an overview cf. Chalmers 2002.

³ I assume throughout that an intension is a content that provides an extension for any suitable possible world and hence can be represented formally by a function from worlds to extensions. Intensions are commonly identified with functions, not senses. I will not follow that custom, since it means confusing senses with their set-theoretic representations. — I stick to the labels for the two intensions used in Chalmers (1996). For a survey of the terminology cf. Chalmers 2002, 6.

⁴ I will by and large follow Chalmers' characterization of neo-descriptivism. It appears that Jackson has recently come around to share Chalmers' views, cf. Chalmers & Jackson 2001, §3.

⁵ One might very well question the viability of this way to characterise epistemic possibilities. I will not do this here, though. However, some of my misgivings show in section 3.1.

⁶ A straightforward two-dimensionalisation of a Kripke-Putnam semantics would fit this description. On such an account, the primary intension of 'water' could be given thus: 'water' applies to x in some context c iff x is of the same kind as the stuff we used to introduce that term in our in fact actual world. The term 'water' would thus be doubly rigid, denoting H₂O in any counterfactual world as well as in any context. We would not know this *a priori*, of course. (If you want to endorse this semantics, you have to adjust the core thesis accordingly.)

⁷ This is *not* simply meant to say that the language we speak if the XYZ-world is the actual world is not the language we actually speak. Neo-descriptivists believe that they can admit a change in language.

⁸ If you are unhappy with the sketched meta-semantic way to account for our intuitions — maybe you object that if I tell the story this way, I assume that the XYZ-world considered as actual contains tokens of 'water' — please note that I will in section 3.3 sketch another way to account for them that is not bound to this assumption.

⁹ As I use it, 'use' captures the narrow linguistic behaviour attached to a certain expression, i.e. the use considered in abstraction from the environment it takes place in.

¹⁰ The general thrust of this argument is well-known. Cf. Devitt/Sterelny 1987, 46ff and Jackson 1998c, 208ff. However, please note that I am concerned with *primary* intensions rather than with common descriptions. This makes things rather different.

¹¹ Cf. Chalmers 2002, 20: "A sentence type D is *a priori* when it is possible for a token of S to be epistemically necessary." Given Chalmers' use of 'token' to characterise variability (cf. Chalmers forthcoming b, 30), this might be understood to pursue the option just outlined.

¹² This is how Chalmers and Jackson spell out their thesis that there is an *a priori* entailment from microphysical(-cum-indexical-cum-phenomenal) truths to ordinary macrophysical truths.

¹³ Maybe you want to hold that ‘primary intension of term *t*’ is by definition what our intuitions concerning the application of *t* across worlds considered as actual add up to. In this case, please understand me to argue that the primary intensions our intuitions add up to are not robust enough to be credible semantic values.

¹⁴ The obvious exception to this rule might be the property of being a liquid.

¹⁵ More precisely, for any set *S* in the disjunction you will find at least one set *S** in the disjunction that contains almost none of the features in *S*.

¹⁶ I assume that features are non-disjunctive.

¹⁷ As the debate sparked by Gettier has demonstrated, it is not too hard to devise cases where we are sure that ‘knowledge’ does not apply.

¹⁸ Actually, I do understand Chalmers’ and Jackson’s talk of worlds ‘not too far from home’ and of ‘presuppositions’ to mean precisely this.

¹⁹ This appears to have been Jackson’s initial position, cf. Jackson 1998, ch. 2. It seems that Jackson has recently dropped this idea. Folk-theories do not even get mentioned in Chalmers & Jackson 2001.

²⁰ I hold that this explanation neither pre-empts nor conflicts with the meta-semantic explanation given in section 3.1.

²¹ These procedures of course are relatives of Putnamian stereotypes, cf. Putnam 1975, 270.

²² As far as I can see, neither Chalmers nor Jackson takes a stance on whether we can, as it were, travel from one world considered as actual to another.

²³ It doesn’t cut any ice with this argument to emphasise that our language does contain lazy occasion terms. For please note that the most obvious lazy occasion terms our language comprises are expressions such as ‘our actual world’ that are tailor-made for counterfactual discourse. This evidently is not true for natural kind terms.

²⁴ I would like to thank Ansgar Beckermann, Michael Schütte, Matthias Adam and a not-so-anonymous referee for comments on earlier versions of this paper.

References

- Almog, J. et. al. (eds.)(1989): Themes From Kaplan, Oxford: Oxford University Press.
- Block, N. and Stalnaker, R. (1999): Conceptual Analysis, Dualism, and the Explanatory Gap, *Philosophical Review* 108, 1-46.
- Byrne, A. (1999): Cosmic Hermeneutics, *Philosophical Perspectives* 13, 347-383.
- Chalmers, D. (1995): The Components of Content
<www.consc.net/papers/content.html>
- Chalmers, D. (1996): The Conscious Mind. In Search of a Fundamental Theory, New York/Oxford: Oxford University Press.
- Chalmers, D./Jackson, F. (2001): Conceptual Analysis and Reductive Explanation, in: *Philosophical Review* 110, 315-360
- Chalmers, D. (2002): The Foundations of Two-Dimensional Semantics, forthcoming in: García-Carpintero, M./Macia, J. (2002): Two-Dimensionalism. European Review of Philosophy <www.consc.net/papers/foundations.html>
- Chalmers, D. (forthcoming a): The Nature of Epistemic Space
<www.consc.net/papers/espace.html>
- Chalmers, D. (forthcoming b): On Sense and Intension
<consc.net/papers/intension.html>
- Davies, M. and Humberstone, L. (1980): Two Notions of Necessity, *Philosophical Studies* 38, 1-30.
- Devitt, M. and Sterelny, K. (1987): Language and Reality, Oxford: Blackwell.
- Fodor, J. (1987): Psychosemantics, Cambridge (Mass.): MIT Press.
- Gleeson, A. (1999): Deducing the Mind, *Inquiry* 42, 385-410.
- Haas-Spohn, U. (1997): The Context-Dependency of Natural Kind Terms, in: W. Kühne et. al (ed.)(1997): Direct Reference, Indexicality, and Propositionale Attitudes, Stanford: CSLI Publications, 333-349.
- Jackson, F. (1994): Armchair Metaphysics, in: Jackson 1998b, 154-176.
- Jackson, F. (1994b): Finding Mind in a Natural World, in: R. Casati et. al. (ed.)(1994): Philosophy and the Cognitive Sciences, Proceedings of the 16th International Wittgenstein Symposium, Vienna: Hölder-Pichler-Tempsky, 101-112.
- Jackson, F. (1994c): Metaphysics by Possible Cases, in: Jackson 1998b, 133-153.
- Jackson, F. (1998): From Metaphysics to Ethics. A Defence of Conceptual Analysis, Oxford: Blackwell.
- Jackson, F. (1998b): Mind, Method and Conditionals. Selected Essays, London/New York: Routledge.
- Jackson, F. (1998c): Reference and Description Revisited, *Philosophical Perspectives* 12, 201-218.
- Jackson, F. (2001): Précis of *From Metaphysics to Ethics*, in: Philosophy and Phenomenological Research LXII, 617-643.

- Kaplan, D. (1977): Demonstratives, in: J. Almog et. al. 1989, 481-563.
- Kaplan, D. (1989): Afterthoughts, in: J. Almog et. al. 1989, 565-614.
- Kripke, S. (1980): Naming and Necessity, Oxford: Blackwell.
- Lewis, D. (1981): Index, Context and Content, in: S. Kanger and S. Öhman (eds.)(1981): *Philosophy and Grammar*, Dordrecht: Reidel 79-100.
- Lewis, D. (1994): David Lewis — Reduction of Mind, in: S. Guttenplan (ed.)(1994): *A Companion to the Philosophy of Mind*, Oxford: Blackwell, 412-431.
- Perry, J. (1977): Frege on Demonstratives, *Philosophical Review* 86, 474-497.
- Perry, J. (1979): The Problem of the Essential Indexical, *Noûs* 13, 5-21.
- Perry, J. (1999): Indexicals and Demonstratives, in: B. Hale and C. Wright (eds.)(1999): *A Companion to the Philosophy of Language*, Oxford: Blackwell, 586-611.
- Putnam, H. (1975): The Meaning of 'Meaning', in: H. Putnam 1975b: *Mind, Language, and Reality*, *Philosophical Papers*, Vol. 2. Cambridge, 215-271.
- Stalnaker, R. (1978): Assertion, *Syntax and Semantics* 9, 315-332.
- Strawson, P.F. (1950): On Referring, *Mind* 59, 320-344.
- Tichy, P. (1983): Kripke on Necessity A Posteriori, *Philosophical Studies* 43, 225-241.
- Yablo, S. (1999): Red, Bitter, Best. Critical Notice of Frank Jackson *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. <http://www.mit.edu/~yablo/rbb.html>; as of January 26, 2002.
- Yablo, S. (2000): Textbook Kripkeanism and the Open Texture of Concepts, *Pacific Philosophical Quarterly* 81, 98-122.
- Yourgrau, P. (ed.)(1990): *Demonstratives*. Oxford Readings in Philosophy, Oxford: Oxford University Press.