



Apriority and Colour Inclusion

Stefan Reining

ADVERTIMENT. La consulta d'aquesta tesi queda condicionada a l'acceptació de les següents condicions d'ús: La difusió d'aquesta tesi per mitjà del servei TDX (www.tdx.cat) i a través del Dipòsit Digital de la UB (diposit.ub.edu) ha estat autoritzada pels titulars dels drets de propietat intel·lectual únicament per a usos privats emmarcats en activitats d'investigació i docència. No s'autoritza la seva reproducció amb finalitats de lucre ni la seva difusió i posada a disposició des d'un lloc aliè al servei TDX ni al Dipòsit Digital de la UB. No s'autoritza la presentació del seu contingut en una finestra o marc aliè a TDX o al Dipòsit Digital de la UB (framing). Aquesta reserva de drets afecta tant al resum de presentació de la tesi com als seus continguts. En la utilització o cita de parts de la tesi és obligat indicar el nom de la persona autora.

ADVERTENCIA. La consulta de esta tesis queda condicionada a la aceptación de las siguientes condiciones de uso: La difusión de esta tesis por medio del servicio TDR (www.tdx.cat) y a través del Repositorio Digital de la UB (diposit.ub.edu) ha sido autorizada por los titulares de los derechos de propiedad intelectual únicamente para usos privados enmarcados en actividades de investigación y docencia. No se autoriza su reproducción con finalidades de lucro ni su difusión y puesta a disposición desde un sitio ajeno al servicio TDR o al Repositorio Digital de la UB. No se autoriza la presentación de su contenido en una ventana o marco ajeno a TDR o al Repositorio Digital de la UB (framing). Esta reserva de derechos afecta tanto al resumen de presentación de la tesis como a sus contenidos. En la utilización o cita de partes de la tesis es obligado indicar el nombre de la persona autora.

WARNING. On having consulted this thesis you're accepting the following use conditions: Spreading this thesis by the TDX (www.tdx.cat) service and by the UB Digital Repository (diposit.ub.edu) has been authorized by the titular of the intellectual property rights only for private uses placed in investigation and teaching activities. Reproduction with lucrative aims is not authorized nor its spreading and availability from a site foreign to the TDX service or to the UB Digital Repository. Introducing its content in a window or frame foreign to the TDX service or to the UB Digital Repository is not authorized (framing). Those rights affect to the presentation summary of the thesis as well as to its contents. In the using or citation of parts of the thesis it's obliged to indicate the name of the author.

University of Barcelona

Faculty of Philosophy

Ph.D. program: Analytic Philosophy

Department of Logic, History and Philosophy of Science

APRIORITY AND COLOUR INCLUSION

Stefan Reining

Supervisor:

Genoveva Martí

May 2014

TABLE OF CONTENTS

Summary	v
Acknowledgements	vii
0. Introduction	1
0.1. The central thesis.....	1
0.2. Some basic assumptions.....	2
0.3. The structure of the dissertation.....	6
1. ‘All crimson things are red’ – Williamson’s argument from skills of imagination and the argument from analogy	9
1.1. The argument from skills of imagination.....	10
1.2. Why the argument fails.....	13
1.3. The role of experience in Norman’s knowledge of R1.....	20
1.4. The argument from analogy.....	24
1.5. Summary and outlook.....	30
2. Examining the argument from analogy, part I: degrees of understanding	33
2.1. Blocking the argument to C1.....	33
2.2. The danger of overgeneralization.....	40
2.3. Grounding the distinction between primary and secondary ways of understanding.....	44
2.4. Why the strategy ultimately fails.....	53
2.5. Summary and outlook.....	55
3. Examining the argument from analogy, part II: P1 and the primary objects of truth	57
3.1. The argument against P1.....	57
3.2. The coarse-grained view and some initial considerations in favour of it.....	61

3.3. Translation.....	67
3.4. Kripke's Paderewski case.....	71
3.5. Conclusions regarding the truth of P1.....	81
3.6. Summary and outlook.....	85
4. Examining the argument from analogy, part III: defending P3.....	87
4.1. Searching for the principle behind P3.....	87
4.2. The criterion for playing a double-role.....	102
4.3. Another argument for P3.....	109
4.4. Summary and outlook.....	114
5. Examining the argument from analogy, part IV: P1 in the light of arguments for colour relativism.....	117
5.1. Colour relativism as a threat to P1.....	117
5.2. The case for colour relativism.....	124
5.3. Some options for resisting the view.....	135
5.4. Resorting to modified versions of P1.....	138
5.5. How standard variation phenomena turn against colour relativism.....	141
5.6. The remaining options.....	147
5.7. Summary.....	153
Conclusions.....	155
Appendix A.....	163
Appendix B.....	165
References.....	167

Summary

My central aim in this dissertation is to propose a new version of local scepticism regarding the a priori, namely, a version of scepticism regarding the apriority of (knowledge of) truths about certain relations between colours. The kind of relation in question is, for instance, expressed by sentences like ‘All ultramarine things are blue’ and ‘Nothing is both red all over and green all over’ – sentences that have, among defenders of the a priori, commonly been regarded as expressing paradigm examples of a priori truths.

In the course of my argumentation for this kind of local scepticism regarding the a priori, I employ a relatively permissive notion of linguistic understanding (inspired by Timothy Williamson’s recent work on the a priori), according to which it is possible to obtain the relevant kind of understanding of colour terms in a certain non-standard way. The relatively permissive notion of linguistic understanding in question is, in turn, based on considerations in favour of a relatively coarse-grained conception of the primary objects of truth. Furthermore, my argumentation for the kind of local scepticism in question is based on considerations in favour of a certain conception of evidentiality, according to which a single experience-token can play both an enabling and an evidential role in the same instance of knowledge, and according to which some of the experience involved in alleged instances of a priori knowledge of the relations among colours in question plays this kind of double-role. Finally, I consider certain empirical phenomena apparently threatening the possibility of coming to understand colour terms in the non-standard way in question, and I argue that the threat posed by these phenomena is more widespread than hitherto acknowledged, and that all available ways of accommodating these phenomena are compatible with my local scepticism regarding the a priori.

Acknowledgements

This dissertation was written in the period from October 2012 to May 2014. The preparatory work on it had started in October 2011. Throughout this time, I received a scholarship from the research network *PERSP – Philosophy of Perspectival Thoughts and Facts*, funded by the Consolider-Ingenio 2010 Scheme (CSD2009-00056) of the Spanish Ministry of Science and Innovation. Without this financial support, my Ph.D. studies at the University of Barcelona would not have been possible. I would therefore like to thank the project coordinator, Manuel García-Carpintero (representatively for all those involved in the decision process), for awarding me this scholarship. I would also like to thank the project administrator, Chiara Panizza, for assisting me in fulfilling the bureaucratic requirements concerning the scholarship.

I received additional financial support from the research group *Reference, Self-Reference and Empirical Data* (FFI2011-25626), funded by the Spanish Ministry of Science and Innovation. I would like to thank the head researcher, José Martínez, for incorporating me into the group.

I would, of course, also like to thank those who directly helped me in elaborating this dissertation. Among those, I would especially like to thank Genoveva Martí for supervising my work, as well as Laurence Bonjour, Yuliya Fadeeva, John Horden, Dan López de Sa, and Timothy Williamson for helpful comments and discussions.

Finally, I would like to thank all those who, in some way or other, intentionally or not, encouraged me in pursuing an academic career. These people – too many to list them here, both from academia as well as from outside the university realm – made it much easier for me to spend most of my energy on my academic work. And without this kind of support, this dissertation would certainly not have come into existence. I hope that those who helped me in providing the necessary amount of dedication and who are themselves pursuing or planning to pursue academic projects will be able to do this under circumstances as comfortable as those that I found myself in during the previous years.

0. Introduction

In this short introductory chapter, I will, first, introduce my central claim, second, state some basic assumptions underlying my argument for my central claim, and third, give a brief preview of the role played by each of the following chapters within my argumentation.

0.1. The central thesis

Among alleged paradigm examples of the a priori, one can often encounter knowledge of (or justification of beliefs in) certain relations between colours. Hilary Putnam's 1956 paper "Reds, Greens, and Logical Analysis", for instance, triggered a debate about whether the apriority of the truth expressed by 'Nothing is red (all over) and green (all over) at the same time' can be explained on the basis of its analyticity – a debate in which only the analyticity, but not the apriority, of this truth is questioned. The same characterization can be found in the writings of the most prominent contemporary defenders of the a priori: Laurence Bonjour, for instance, lists the proposition that no surface can be uniformly red and uniformly blue at the same time as one of five supposedly clear examples of propositions that can be known a priori (cf. Bonjour 2005, p. 100; see also Bonjour 1998, pp. 100-102). And Christopher Peacocke holds that it "seems that no more than grasp of the relevant colour concepts is required for one to be in a position to appreciate the incompatibility of a surface being wholly definitely red and wholly definitely green" (2000, p. 256; cf. also Audi 1998, p. 101), which is supposed to imply the apriority of the corresponding truth.

My aim in this dissertation will be to challenge this common way of characterizing the kind of relation in question. More precisely, my aim will be to argue that what I call 'colour inclusion sentences' do not express a priori truths. The notion of a colour inclusion sentence is supposed to cover all non-logically-true sentences expressing a relation of inclusion, or exclusion, of colour properties possessed by a given object at a given time. That is, the notion of a colour inclusion sentence is supposed to cover sentences like 'All ultramarine things are blue', as well as 'Nothing is both blue all over and red all over at the same time'; but it is not supposed to cover sentences like 'All

dark red things are red', since the latter is a logical truth (to the extent that sentences are truths at all).

My argument for the claim that colour inclusion sentences do not express a priori truths will still be compatible with the apriority of other paradigm instances of the a priori, such as logical and arithmetic truths, thereby merely constituting a local form of scepticism regarding the a priori.

0.2. Some basic assumptions

In my argumentation against the apriority of what is expressed by colour inclusion sentences, I will not assume any specific conception of the a priori. Otherwise, while my argumentation might be interesting for those sharing the respective conception of the a priori, it would not be of much interest for those endorsing a contrary conception. Rather, I will merely assume a relatively uncontroversial necessary condition – namely, I will assume that for an instance of knowledge to be a priori, experience (that is, perception or introspection) must play no evidential role in it. This assumption is widely shared among contemporary defenders of the existence of a priori knowledge, both by proponents of purely negative conceptions of the a priori and by proponents of more demanding positive conceptions. That is, the assumption that for an instance of knowledge to be a priori, experience must play no evidential role in it is, for instance, shared by Casullo 2003 in his purely negative conception of a priori knowledge, according to which “a priori knowledge is knowledge whose justification is nonexperiential“ (p. 81) – a conception which is simply exhausted by the assumption in question, given that the terms ‘evidential role’ and ‘justificatory role’ are used interchangeably in the relevant debate. In a recent defense of the a priori against various naturalistic challenges, Jenkins 2013 conceives of a priori knowledge or justification as knowledge or justification which is “epistemically independent of empirical evidence” (p. 275) – a negative conception I take to be equivalent (to any relevant extent) to both Casullo’s conception and my above-mentioned assumption.¹ Furthermore, the

¹ Kant’s remarks at the beginning of the introduction to the second edition of his *Critique of Pure Reason* may suggest that his conception of a priori knowledge does not differ from the negative conception just mentioned. That is, his decision to call exactly those instances of knowledge ‘a priori’ that “occur *absolutely* independently of all experience” (Kant 1998, p. 137), together with his claim that “all our cognition commences *with* experience, yet it does not on that account all arise *from* experience” (ibid., p.

assumption in question is, for instance, shared by Bonjour 1998 in his more demanding positive conception of a priori knowledge, according to which the only evidential basis of a priori knowledge is an *intellectual* “seeing” of the necessity of the respective proposition (cf. p. 106). Of course, the assumption in question might be fleshed out in different ways, given different possible conceptions of an evidential role, that is, given different possible conceptions of what it means for some experience to play an evidential (rather than a merely enabling) role in a given instance of knowledge. At this early stage, I will not put forward any assumption on what it means for some experience to play an evidential role. Rather, an assumption of this kind will constitute one of the premises of my central argument, and it will be challenged in the course of examining the argument’s soundness.

It should be noted that my assumption that for an instance of knowledge to be a priori, experience must play no evidential role in it is not supposed to carry with it the assumption that a priori knowledge is empirically indefeasible (that is, the assumption that a priori knowledge could not be undermined by the occurrence of empirical evidence). The latter assumption is commonly regarded as incompatible with the apriority of some of the supposed paradigm cases of a priori truth, such as the truth of arithmetic equations. As noted by Field 2000, p. 118, it seems that empirical indefeasibility, understood as entailing that there cannot be empirical evidence against the respective truth, thereby entails that there can also be no empirical evidence *for* the respective truth. And it is commonly assumed that this latter requirement does not hold regarding some supposed paradigm cases of a priori truth (see, for instance, Kripke 1980, p. 35, and Kitcher 1984, p. 22; see, however, Field 2000, p. 118, for a possible way of denying this assumption). Thus, should it (surprisingly) turn out that all instances of knowledge in which experience does not play an evidential role are necessarily empirically indefeasible – that is, if it, for instance, (surprisingly) turned out that (experience of) the mere absence of possible empirical defeaters automatically plays an evidential role in all instances of empirically defeasible knowledge (see Summerfield 1991, pp. 41ff., for a rejection of this claim on the basis of the notions of positive and negative epistemic dependence) –, then defenders of the a priori would

136), may suggest that he calls exactly those instances of knowledge ‘a priori’ in which experience does not play an evidential, but merely a purely enabling, role. However, interpreting the second of the two passages just cited is certainly a delicate issue; and I will not engage in interpreting historical passages by Kant or others, given that it will not be essential for the purposes of this dissertation.

have reason to be concerned. In my argumentation, I will stay neutral regarding this issue.

As already indicated by some of the above formulations (especially by the formulation of my central claim that colour inclusion sentences do not express a priori truths), I will further assume that apriority can felicitously be ascribed to the objects of truth. However, I do not thereby assume that the objects of truth are the *primary* objects of apriority, or that apriority is a non-epistemic property. Rather, I assume that a truth is a priori if and only if there could be a priori knowledge of it (that is, if and only if in some possible worlds, the truth in question is known a priori), and I am ready to accept that the notion of a priori truth is a merely derivative one. Moreover, I will stay neutral regarding whether the notion of a priori knowledge is itself also a derivative one – derivative, for instance, on the notion of a priori warrant or a priori (doxastic) justification (see, e.g., Cassam 2000: “For one's knowledge of *p* to count as a priori, one's justification for believing this proposition must be a priori” (p. 44)). Those regarding warrant or justification as the primary object of apriority are faced with the further issue of whether one could be a priori warranted or justified in believing a falsehood (see, e.g., Hale 1987 for a positive answer: “But why assume, in the first place, that a priori warrants *must* produce true beliefs?” (p. 129)). Since I will stay neutral regarding whether the notions of a priori truth and a priori knowledge are derivative on the notion of a priori warrant or justification, I will also stay neutral regarding the issue just mentioned.

Assuming that apriority can felicitously be ascribed to the objects of truth, and that a truth is a priori if and only if it is possible to have a priori knowledge of it, evokes the following questions (posed by Kripke 1980, pp. 34f.): “[P]ossible for whom? For God? For the Martians? Or just for people with minds like ours?” If some God-like being, existing in some possible worlds, were to possess a priori knowledge of all sorts of truths that could only be known a posteriori by human beings, then a non-relativized notion of a priori truth would simply cover all truths, and should therefore be regarded as uninteresting. And as Kripke notes (cf. p. 35), resorting to a conception of a priori truth according to which a truth is a priori if and only if it is *not* possible to have a *posteriori* knowledge of it would be of no help, since it is quite plausible to suppose that paradigm cases of truths that can be known a priori, such as truths of logic and arithmetic, can as well be known a posteriori, e.g., via relying on testimony or a

calculator. Thus, one may feel tempted to share Kripke's suspicion regarding the usefulness of the notion of a priori truth:

"It might be best therefore, instead of using the phrase 'a priori truth', to the extent that one uses it at all, to stick to the question of whether a particular person or knower knows something a priori or believes it true on the basis of a priori evidence." (1980, p. 35)

However, concerning the questions posed by Kripke that I have cited above, I tend to assume that no possible being whatsoever, regardless of the extent to which it deserves the label 'God-like', is able to acquire a priori knowledge of truths that could not in principle also be known a priori by human beings. That is, the kind of being in question may be able to acquire a priori knowledge of truths that, due to mere limits of time and cognitive capacity, no actual human being can acquire (such as a priori knowledge of the truth of an extremely long arithmetic equation). However, when it comes to truths that, independently of any extension of time and cognitive capacities, human beings could in principle only come to know by relying on empirical evidence, I feel tempted to assume that even the divine being in question is not able to acquire knowledge that deserves the label 'a priori'. That is, whatever the connection between the being's mind and the respective empirical fact that endows the being with knowledge of this fact consists in, the connection has to be experiential in some way or other (though perhaps not in any *familiar* way). A being possessing true beliefs about so-called empirical facts without any (direct or indirect) causal connection involving experience (understood in a wide sense) should, so I tend to assume, be regarded as getting things right merely by luck, and thus as not possessing the respective knowledge, neither a priori knowledge nor a posteriori knowledge. It is quite plausible to suppose that there are possible worlds in which human beings, with limited cognitive capacities, are in the kind of lucky position just described with respect to a large set of empirical facts. And it is also plausible to suppose that these lucky human beings should not be regarded as knowing the respective truths. Any divine being, with unlimited cognitive and physical capacities, who gets things right without any causal connection involving experience (understood in a wide sense) should, so I tend to assume, be regarded as being in the same kind of lucky position as the human beings just described, and therefore as not knowing the respective truths. However, such an assumption may well be regarded as contentious. And making this assumption is, moreover, not at all necessary for my purposes. Therefore, I will simply leave open whether the notion of a priori truth that I

employ is subject-relative or not. My considerations regarding the apriority of what is expressed by colour inclusion sentences exclusively concern *human* knowledge of these truths. To the extent that I am right regarding the issue of divine beings described above, the notion of a priori truth employed here is supposed to be a subject-independent one. Should it turn out that I am wrong regarding the issue in question (such that my central claim that colour inclusion sentences do not express a priori truths would be false, given a subject-independent notion of a priori truth), I could still resort to a species-relative notion of a priori truth, since my considerations in favour of my central claim do not hinge on any contingent limitations of the cognitive capacities of human beings.

Finally, it should be pointed out that my assumption that apriority can felicitously be ascribed to the objects of truth is not supposed to carry with it any assumption regarding the nature of the primary objects of truth, that is, regarding whether the primary objects of truth are propositions, sentences, etc. At some point during the defense of my central argument, I will argue for a certain claim regarding the fineness of grain of the primary objects of truth. However, even at that point, the question whether the primary objects of truth are propositions, sentences, etc. will be left open.

0.3. The structure of the dissertation

My case for my central thesis that colour inclusion sentences do not express a priori truths mainly consists in one central argument, which I will call ‘the argument from analogy’. This argument, together with some short initial considerations in favour of its three premises, will be presented toward the end of chapter one. The argument evolves out of a discussion of Timothy Williamson’s recent treatment of a supposed paradigm instance of a priori knowledge of what is expressed by a certain colour inclusion sentence. Williamson 2013 argues that, in the instance of knowledge in question, as well as in a certain supposed paradigm instance of a posteriori knowledge, experience neither plays a purely enabling nor a strictly evidential role – a result calling into question the discriminatory significance of the labels ‘a priori’ and ‘a posteriori’. In chapter 1, I first argue that Williamson is wrong in claiming that in the supposed paradigm instance of a posteriori knowledge in question, experience does not play a strictly evidential role. The argument from analogy, however, which will be presented

thereafter, would, if successful, vindicate Williamson's claim that there is no epistemologically significant difference between the supposed paradigm instance of a priori knowledge and the supposed paradigm instance of a posteriori knowledge in question. That is, the argument from analogy is an argument to the conclusion that experience does not only play a strictly evidential role in the supposed paradigm instance of a posteriori knowledge in question, but also in the supposed paradigm instance of a priori knowledge in question (namely, in the supposed paradigm instance of a priori knowledge of what is expressed by a certain colour inclusion sentence).

Together with certain plausible generalizing assumptions, the soundness of the argument for analogy would constitute a case for my central thesis. Therefore, I will devote the whole remainder of the dissertation (chapters 2-5) to a defense of the argument's soundness. This defense will partly consist in rejections of objections against the truth of some of the argument's premises (especially in chapters 2 and 5), but it will also partly consist in positive considerations in favour of the truth of some of the argument's premises (especially in chapters 3 and 4). In chapter 5, I will (among other things) consider certain reasons for claiming that colour inclusion sentences do not express truths (or falsehoods) at all – a claim incompatible with the soundness of the argument from analogy, but also incompatible with the falsity of my central thesis.

1. ‘All crimson things are red’ – Williamson’s argument from skills of imagination and the argument from analogy

My central aim in this first chapter is to present the outlines of an argument to the conclusion that paradigm ways of knowing truths expressed by colour inclusion sentences are not a priori. In particular, the argument to be presented is directed against the apriority of paradigm ways of knowing that all crimson things are red – an (exchangeable) example of colour inclusion taken from Timothy Williamson’s recent paper “How Deep is the Distinction between A Priori and A Posteriori Knowledge?”. In this paper, Williamson argues for the claim that the distinction between the a priori and the a posteriori is merely a superficial one – one that “does not cut at the epistemological joints” (2013, p. 294). The central part of his argumentative strategy consists in “comparing what would usually be regarded as a clear case of a priori knowledge with what would usually be regarded as a clear case of a posteriori knowledge” (ibid.), and attempting to show that there is no deep epistemological difference between the two kinds of cases.

In section 1.1, I will reiterate Williamson’s argument for the claim that there is no deep epistemological difference between the two examples he considers. I will call this argument the ‘argument from skills of imagination’. In section 1.2, I will show that the latter is unsound. In section 1.3, I will present an argument for the claim that there *is* a deep epistemological difference between the two cases in question. I will then highlight the main shortcoming of this type of argument. Finally, in section 1.4, I will present an argumentative strategy, partly inspired by Williamson’s work on conditions for linguistic understanding, that might be used to vindicate Williamson’s claim that there is no deep epistemological difference between the two cases in question. That is, I will present an argument to the conclusion that experience plays an evidential role in paradigmatic cases of knowing that all crimson things are red. This argument, which will be called the ‘argument from analogy’, is generalizable to all other truths expressed by colour inclusion sentences, and its function will not only be to vindicate Williamson’s claim that there is no deep epistemological difference between the two cases in question, but also to constitute the main part of my case for my central thesis that colour inclusion sentences do not express a priori truths. In later chapters, the argument’s soundness will be examined in more detail.

1.1. The argument from skills of imagination

In his argument for the claim that the distinction between the a priori and the a posteriori is merely a superficial one, Williamson compares paradigmatic instances of knowledge of the supposedly a priori truth

(R1) All crimson things are red.²

with paradigmatic instances of knowledge of the supposedly a posteriori truth

(R2) All recent volumes of *Who's Who* are red.

He begins his argumentation by construing the following scenario of someone coming to know R1 (cf. Williamson 2013, p. 295). This person, Norman, has never been explicitly taught about any co-occurrence relation between the two colours. Instead, he has acquired the two colour concepts independently of one another, through ostension. He then gets confronted with the question whether R1 is true, and he comes to know the answer to this question not by looking at, or remembering the colour of, any particular object, but merely by using “his skill in making visual judgments with ‘crimson’ to visually imagine a sample of crimson” and by using “his skill in making visual judgments with ‘red’ to judge, within the imaginative supposition, ‘It is red’” (ibid.).

Williamson admits that this description of a cognitive process would have to be enriched in order to guarantee that the process in question yields knowledge of R1. That is, Norman’s skill in making visual judgments with ‘crimson’ should not only enable him to reliably imagine some arbitrary sample of crimson, but it should, in the case described, enable him to reliably imagine a *central prototype* of crimson (cf. Williamson 2013, p. 296); and the same should hold regarding his skill in making visual judgments with ‘red’. Moreover, his skills in making visual judgments with ‘crimson’ and ‘red’ should enable him to reliably assess how much a colour shade may differ from the imagined central prototypes of crimson and red so as to still count as a shade of crimson or red, respectively. This kind of assessment skill is necessary in order to

² For present purposes, nothing hinges on whether R1 should, as Williamson assumes, really be regarded as a truth, or whether crimson should instead be regarded as a borderline case of red. Nothing essential would be lost if the sentence ‘All crimson things are red’ were replaced by ‘All scarlet things are red’, the latter expressing a clearer case of colour inclusion.

differentiate between the true claim that all crimson things are red and the false claim that all red things are crimson (cf. *ibid.*). But even though these cognitive processes could in principle be spelled out in detail, such an endeavor is, as Williamson remarks, not necessary for the purposes of the comparison (cf. *ibid.*).

Regarding the acquisition of knowledge of R2, the person to be imagined is again Norman. In addition to what has been stipulated about him in the first story, it is further to be imagined (cf. Williamson 2013, p. 296) that Norman has never previously been confronted with the question whether R2 is true. He is a competent user of all the single expressions contained in R2, and he is reliable in judging whether something is a recent volume of *Who's Who*, and whether something is red. According to Williamson, these competences and faculties enable Norman to come to know R2,

“without looking at any recent volumes of *Who's Who* to check whether they are red, or even remembering any recent volumes of *Who's Who* to check whether they were red, or any other new exercise of perception or memory” (*ibid.*).

Instead, so Williamson claims, Norman can come to know R2 analogously to the way in which he comes to know R1, that is, by using his skill in making visual judgments with ‘recent volume of *Who's Who*’ to imagine a (central prototype of a) recent volume of *Who's Who*, and by using “his skill in making visual judgments with ‘red’ to judge, within the imaginative supposition, ‘It is red’” (*ibid.*). And in order to make the crucial point unmistakably clear, Williamson repeats: “No episodic memories of prior experiences, for example of recent volumes of *Who's Who*, play any role” (*ibid.*).

He then draws the following conclusion from the comparison. Norman’s knowledge of R1 and R2 are clear cases of what would generally be regarded as instances of a priori and a posteriori knowledge, respectively. But the cognitive skills underlying the two processes of knowledge acquisition are “almost exactly similar.” (Williamson 2013, p. 296) This, Williamson suggests, shows that there is no deep epistemological difference between R1 and R2, which, in turn, suggests that “the a priori – a posteriori distinction is epistemologically shallow” (*ibid.*, p. 297). Of course, R1 and R2 differ in modal status, since the truth of R1 is necessary, while the truth of R2 is contingent. But this is of no *epistemological* significance, since “the distribution of errors in modal space may be much the same in the two cases” (*ibid.*, p. 299). That is, even though R1 is true in all possible worlds, while R2 is true in some and false in others, one may well be as error-prone regarding the truth or falsity of R1 as regarding

the truth or falsity of R2. Therefore, Williamson claims that “[t]he main effect of the modal difference between R1 and R2 may be to distract us from the epistemological similarity” (ibid., p. 300).

As already stated in the introductory chapter, I will make the relatively common assumption that in order for an instance of knowledge to count as a priori, experience must play no evidential role in it. Obviously, the term ‘evidential’ is in need of clarification, as well as the complementary term ‘enabling’. A usual way of defining the latter is to state that experience is enabling in so far as it contributes to the understanding of the relevant expressions (or to the grasp of the corresponding truth). For example, knowledge of the fact that all bachelors are unmarried certainly requires some kind of experience. But, according to what is commonly assumed, this kind of experience is only necessary for the kind of knowledge in question in so far as it provides for the subject’s understanding of the sentence ‘All bachelors are unmarried’ (or for the subject’s grasp of what is expressed by this sentence). In this specific example, the kind of experience in question includes all kinds of perceptions that contribute to the subject’s understanding of the sentence (or grasp of the corresponding truth), as well as the kind of sensory input that was necessary for the subject to develop the general faculty of making judgments. In contrast, experience is usually regarded as playing an evidential role in so far as the role it plays is not purely enabling. For example, the role that experience plays in the knowledge of the fact that it rained yesterday is, according to what is commonly assumed, not purely enabling, since the mere understanding of the sentence in question (or grasp of the corresponding truth) has to be combined with some further experience in order to put the subject in a position to know that it rained yesterday.

Regarding the two cases of knowledge that Williamson considers, he claims that the distinction between an enabling and an evidential role of experience is of no use for demonstrating a deep epistemological difference, because “the role of experience in both cases is more than purely enabling but less than strictly evidential.” (Williamson 2013, p. 298) That is, he neither regards the role of experience played in Norman’s knowledge of R1 as purely enabling, nor does he regard the role of experience played in Norman’s knowledge of R2 as strictly evidential.

In order to show that the role played by experience in Norman’s knowledge of R1 is more than purely enabling, Williamson construes the case of Norbert (cf. Williamson 2013, p. 297), who, like Norman, has learned to use the terms ‘crimson’ and ‘red’ so as

to be reliable in correctly applying these terms to objects in his surroundings. However, Norbert has had less experience than Norman in using the term ‘crimson’, which “makes him less skilful than Norman in imagining a crimson sample” (ibid.). Despite this lack of competence concerning the faculty of offline processing, Williamson claims that “[b]y normal standards [Norbert] is linguistically competent with both words [and] grasps proposition [R1]” (ibid.). However, due to Norbert’s lack of skill in imagining a sample of crimson, he does not come to know R1. Thus, according to Williamson, the role played by experience in Norman’s knowledge of R1 is more than purely enabling, since, as the case of Norbert is supposed to illustrate, knowledge of R1 requires more from experience than merely enabling one to understand the respective sentence or grasp the respective truth.

Regarding Norman’s knowledge of R2, Williamson claims that “[t]he only residue of his experience of recent volumes of *Who’s Who* active in his knowledge of [R2] is his skill in recognizing and imagining such volumes” (Williamson 2013, p. 298), and he regards this role of experience as “less than strictly evidential” (ibid.). He therefore concludes that no deep epistemological distinction between Norman’s knowledge of R1 and of R2 can be drawn by appealing to the conceptual difference between an enabling and an evidential role of experience.

1.2. Why the argument fails

Williamson devotes some significant space (cf. Williamson 2013, pp. 301ff.) to arguing that Norman’s way of coming to know R1 is not relevantly different from people’s supposedly a priori knowledge of mathematical and logical truths. In what follows, I will not contest this claim.³ Instead, I will focus on Williamson’s claim that the role played by experience in Norman’s knowledge of R1 is not significantly different from the role played by experience in Norman’s knowledge of R2. As described in the previous section, the latter is based on considerations in favour of the claim that the role played by experience in both of these cases of knowledge is neither purely enabling nor strictly evidential. I will argue that the reasons that Williamson offers for accepting this

³ For objections against the claim that Norman’s knowledge of R1 is not relevantly different from supposedly a priori instances of knowledge of mathematical truths (in particular, of the power set axiom), see Jenkins & Kasaki (forthcoming), section 3.

latter claim are not convincing. More precisely, I will argue that his considerations regarding skills of imagination fail to provide a good reason for accepting the claim that the role played by experience in Norman's knowledge of R2 is not a strictly evidential one. I thereby argue that Williamson's argument for the claim that there is no significant epistemological difference between Norman's knowledge of R1 and Norman's knowledge of R2 is unsound.

In order to see that the role played by experience in Norman's knowledge of R2 should be regarded as strictly⁴ evidential, it is necessary to have a closer look at Williamson's claim that "[t]he only residue of [Norman's] experience of recent volumes of *Who's Who* active in his knowledge of [R2] is his skill in recognizing and imagining such volumes." (Williamson 2013, p. 298) I grant that this claim is true. However, I claim that the skill of imagining involved in Norman's knowledge of R2 is epistemically more demanding than it might superficially appear. That is, I will argue that the kind of imagining that figures in Norman's knowledge of R2 requires experience to play a role that should most plausibly be regarded as an evidential one.

The kind of imagining that figures in Norman's knowledge of R2 involves the skill of imagining recent volumes of *Who's Who*. This formulation of the skill is in need of clarification, for it can be used to refer to quite different skills. There are mainly two possible readings. According to one of these two readings, the skill of imagining recent volumes of *Who's Who* consists in the following:

- a) the skill of visually representing to oneself some possible thing that *could* be a recent volume of *Who's Who* (while judging that it could be such a volume).

Put into less inelegant terms, the skill described under a) might be redescribed as the skill of imagining what a recent volume of *Who's Who* might look like. Exercising this skill could, for example, consist in visually representing to oneself some possible green book with the expression "Who's Who" written on the cover, while judging that a recent volume of *Who's Who* might look exactly like that.

According to the second of the two readings, the skill of imagining recent volumes of *Who's Who* consists in the following:

⁴ For the moment, I interpret the expression 'strictly evidential' in a way according to which sentences of the form 'x is strictly evidential' are true iff corresponding sentences of the form 'it is strictly true that x is evidential' are true. From now on, I will only use the shorter expression 'evidential', intending it to be interpreted in this way. However, I will address this issue again below.

- b) the skill of visually representing to oneself a stereotype (regarding colour) of those things that in fact *are* recent volumes of *Who's Who*.

Obviously, what is needed for Norman's knowledge of R2 is not merely the skill described under a), but also the one described under b); for the skill described under a) might at best enable him to know that all recent volumes of *Who's Who* *could* be red.

Now, in order to see that the role of experience played in Norman's knowledge of R2 should be regarded as evidential, one merely has to consider the kind of experience that Norman had to have in order to reliably exercise the skill described under b). In order for Norman to exercise the skill described under b) as reliably as his knowledge of R2 requires, he had to take a (somewhat) recent look at volumes of *Who's Who*. Of course, Norman might in principle also gather the relevant information via testimony, but the example, as stipulated by Williamson, is one in which the relevant information is obtained through direct observation. The reason why Norman's reliable exercise of the skill described under b) requires him to have taken a somewhat recent look at volumes of *Who's Who* is simply that what is a stereotype (regarding colour) of recent volumes of *Who's Who*, relative to some time *t*, might not be a stereotype (regarding colour) of recent volumes of *Who's Who*, relative to some other time. In other words: if the stereotypical colour of recent volumes of *Who's Who* is also the stereotypical colour of past volumes, then this is contingently so. Of course, it might in principle be the case that Norman has not taken a look at volumes of *Who's Who* for a long time, but that he has good reasons to believe that new volumes of *Who's Who* have the same colour as older ones. In such a case, Norman's reliable exercise of the skill described under b) would *not* require him to have taken a somewhat recent look at some of the volumes. But again, the example, as stipulated by Williamson, is not supposed to be read in this way. Therefore, putting it into lax terms, Norman's knowledge of R2 requires him to be visually up-to-date.

However, the claim that Norman's knowledge of R2 requires him to be visually up-to-date might still not be regarded as a good reason for accepting the claim that the role played by experience in Norman's knowledge of R2 is an evidential one. In order to overcome these reservations, it might be helpful to distinguish between two different readings of the claim that Norman's knowledge of R2 requires him to be visually up-to-date. Once it has been clarified which of the two readings is the intended one, the

conclusion that the role played by experience in Norman's knowledge of R2 is evidential should no longer be denied.

The distinction between two different readings of the claim that Norman's knowledge of R2 requires him to be visually up-to-date amounts to a distinction between two different kinds of requirement. The first kind of requirement is one that is merely due to people's limited memory capacities. It might, for example, be that Norman loses his skill of offline representation regarding certain concepts if he does not from time to time refuel his visual memory by direct observation. If the claim that Norman's knowledge of R2 requires him to be visually up-to-date is read in this way, then its truth does indeed not provide a good reason for accepting that the role played by experience in Norman's knowledge of R2 is an evidential one. However, the intended reading is a different one. The kind of requirement corresponding to the intended reading is one that holds *in principle*, regardless of the limitations of people's memory capacities.

In order to illustrate that this latter kind of requirement can only be fulfilled if some of the experience involved is evidential, consider a modified version of Williamson's case which differs from the original version only insofar as Norman has been space-travelling for many years, having had no contact whatsoever with his fellow Earthians. Since the case is construed in such a way that he never had any good reason to suppose that future volumes of *Who's Who* would have the same colour as previous ones, coming to know R2 would require him to travel back to Earth (or to make some other kind of contact) and *check* the colour of recent volumes of *Who's Who*. And this requirement is a principled one: it holds even though – as Williamson would admit – Norman understands all of the expressions involved, and it holds independently of whether his visual memory has to be refueled from time to time.

This case illustrates that a reliable exercise of the skill of imagining involved in Norman's knowledge of R2 is epistemically more demanding than it might superficially appear. In the modified version of the case, a reliable exercise of this skill requires Norman to make some contact with his home planet in order to check the colour of recent volumes of *Who is Who* – even though he grasps R2 perfectly well, and independently of the quality of his visual memory. It would be extremely implausible to claim that the role played by this kind of experience is not an evidential one. In the original version of the case, checking the colour of recent volumes of *Who is Who* does not require as much effort for Norman as it does in the modified version, since in the

original version he has never left Earth. But this makes no difference regarding the general evidential character of the experience required for reliably exercising the skill of imagining described under b). It can be concluded that experience plays an evidential role in Norman's knowledge of R2, and that Williamson's argument for the claim that there is no deep epistemological difference between Norman's knowledge of R1 and his knowledge of R2 is therefore unsound.

An earlier version of Williamson's argument for the claim that the role played by experience in what would be regarded as clear cases of a posteriori knowledge can be less than evidential already appears in *The Philosophy of Philosophy* (cf. Williamson 2007, pp. 165ff.). Although this earlier version is shorter and less elaborated than the one discussed above, it is still worthwhile to have a look at the example he uses there, because this example is not completely analogous to R2, and it might therefore not be completely obvious how to apply the above considerations to it. Here is the example (cf. *ibid.*, p. 167):

(ANT) If two marks had been nine inches apart, they would have been further apart than the front and back legs of an ant.

One apparent disanalogy between R2 and ANT consists in the latter being a counterfactual conditional. (The reason why Williamson chooses this kind of example is that the shorter version of the argument in *The Philosophy of Philosophy* has the function of determining the consequences of the supposed fact that the epistemology of modal truths reduces to the epistemology of counterfactual conditionals; but this larger context is not relevant for present purposes.) Translated into possible worlds talk, ANT says that in the closest possible worlds (to the actual world) in which there are two marks that are nine inches apart, these two marks are further apart than the front and back legs of an ant. But this reformulated version of ANT is still quite unclear. For it can still be interpreted in relevantly different ways, including the following:

(ANTa) In the closest possible worlds (to the actual world) in which there are two marks that are nine inches apart, these two marks are further apart than the front and back legs of *most ants presently existing in the actual world*.

(ANTb) In the closest possible worlds (to the actual world) in which there are two marks that are nine inches apart, these two marks are further apart than the front and back legs of *most ants presently existing in the respective possible worlds*.

Given that it is a contingent matter whether the actual world is one in which there are two marks that are nine inches apart, neither of the two readings entails the other. They thus have to be considered separately. I start with ANTa.⁵

Imagine that Norman knows ANTa, and that, analogously to his knowledge of R1 and R2, his knowledge of ANTa is not based on any testimonial evidence, but instead only on his perceptual acquaintance with ants and with nine-inch distances. In order to see that the role played by experience in Norman's knowledge of ANTa is an evidential one, consider again the scenario in which he has been space-travelling for many years. Williamson himself regards ANT as contingently true (cf. Williamson 2007, p. 167), thereby conceding the possibility of ants whose front and back legs are at least nine inches apart. In order for Norman, in his spaceship, to know ANTa, it is not sufficient for him to be still reliable in accurately measuring, in his imagination, the size that ants had when he was still on Earth. For it might have been that, during his long travel, ants have died out, except for some genetically manipulated ones whose front and back legs are more than nine inches apart. This scenario is, of course, far less likely to occur than one in which recent volumes of *Who's Who* are green instead of red. But nevertheless, a reliable exercise of the skill of imagining the distance between the front and back legs of most ants presently existing in the actual world – which, according to Williamson, would be involved in Norman's coming to know ANTa – would require Norman to check what is going on on his home planet. And, analogously to what has been stated in the discussion of R2 above, it would be extremely implausible to claim that the role played by this kind of experience is not an evidential one. And, also analogously to what has been stated regarding R2, construing the case as one in which Norman is in

⁵ Other interpretational variations arise from replacing the expression 'most ants' in ANTa and ANTb with 'all ants', or from replacing the expression 'in the actual world' in ANTa, or the expression 'in the respective possible worlds' in ANTb, with 'in all possible worlds'. The only important thing to note regarding these possible variations of ANTa and ANTb is that they would make the respective version stronger, that is, the resulting versions of ANTa and ANTb would entail the respective original version, but not vice versa. So if it can be demonstrated that experience must play an evidential role in one's knowledge of ANTa and ANTb, then the same must hold regarding one's knowledge of any of the possible variations of ANTa and ANTb just mentioned. It will therefore suffice to consider ANTa and ANTb.

outer space, rather than one in which he has never left Earth, makes no difference regarding the general evidential character of the kind of experience required for the kind of imagining involved.

Concerning ANT_b, matters are not significantly different. If the actual world is one in which there are two marks that are nine inches apart, the size of ants in the actual world is again relevant. And even if there were no such two marks in the actual world, Norman would, in order to reliably estimate the size of ants in the relevant possible worlds, still have to visit Earth or contact his fellow Earthians. In conclusion, then, replacing R2 with either ANT_a or ANT_b makes no difference regarding the soundness of Williamson's argument.

Above (in fn. 4), I announced that I would use the term 'evidential' as an abbreviation of Williamson's expression 'strictly evidential', and that I would interpret the expression 'strictly evidential' in a way according to which sentences of the form 'x is strictly evidential' are true iff corresponding sentences of the form 'it is strictly true that x is evidential' are true. In order to be fair to Williamson, it should be pointed out that under at least one different, and not too far-fetched, interpretation of the expression 'strictly evidential', my attempt at refuting his claim that experience plays no strictly evidential role in Norman's knowledge of R2 (or of ANT_a or ANT_b) can no longer be regarded as successful. The interpretation I have in mind is one according to which sentences of the form 'memory plays a strictly evidential role in S's knowledge of p' are true only if S's belief that p is based on some *episodic* memory of particular experiences, and not merely, as in the case of Norman's knowledge of R2, on some *generic* memory of what certain objects look like. Regarding Norman's knowledge of R2, Williamson stipulates that "[n]o episodic memories of prior experiences, for example of recent volumes of *Who's Who*, play any role." (2013, p. 296). Thus, under the possible interpretation of the expression 'strictly evidential' mentioned above, it would have to be conceded that experience plays no strictly evidential role in Norman's knowledge of R2.

However, it is simply not plausible to regard experience's playing a strictly evidential role, understood in such a demanding way, as necessary for an instance of knowledge to deserve the label 'a posteriori'. Given the above considerations regarding requirements of being visually up-to-date, a good deal of argument would have to be provided in order to show that the distinction between episodic and generic memory is of any significance here. Without such arguments being provided, I take it to be most

reasonable to assume that the less demanding interpretation of the expression ‘strictly evidential’ that I have used throughout my argumentation in the present section (the one announced in fn. 4) is more adequate regarding the epistemological issues at hand. And as I have argued, under this interpretation of the expression ‘strictly evidential’, Williamson’s claim that experience plays no strictly evidential role in Norman’s knowledge of R2 is false. (Henceforth, I will again go on to abbreviate ‘strictly evidential’ with the shorter expression ‘evidential’.)

1.3. The role of experience in Norman’s knowledge of R1

I have argued that the role played by experience in Norman’s knowledge of R2 is an evidential one. Hence, in order to vindicate Williamson’s claim that there is no significant epistemological difference between Norman’s knowledge of R2 and his knowledge of R1, one would have to show that in Norman’s knowledge of R1 experience also plays an evidential role. As stated above, Williamson claims that the role played by experience in Norman’s knowledge of R1 is *not* an evidential one. In section 1.4, I will present a possible strategy for proving Williamson wrong in this respect. This strategy mainly consists in contrasting his original version of the scenario with a certain modified version of it, and in drawing a certain analogy between the two scenarios. In the present section, however, I will first briefly explain why Williamson’s arguments for denying that experience plays an evidential role in Norman’s knowledge of R1 are unsatisfactory. I will present an argument for the claim that experience plays no evidential role in Norman’s knowledge of R1, but I will also argue that the kind of argument in question is invalid.

In the original version of the scenario, Norman acquires the words ‘crimson’ and ‘red’ by means of learning about the superficial⁶ phenomenal properties of crimson things and red things, respectively. As Williamson puts it,

“[h]e learns ‘crimson’ by being shown samples to which it applies and samples to which it does not apply, and told which are which. He learns ‘red’ in a parallel but causally independent way. [...] Through practice and feedback, he becomes very skilful in judging by eye whether something is crimson, and whether something is red.” (Williamson 2013, p. 295)

⁶ I will use the expressions ‘superficial phenomenal properties’ and ‘superficial looks’ in order to refer to the way things look to the unaided eye, rather than to the way they look when seen through a microscope.

Norman is then asked to make an offline judgment regarding whether all crimson things are red. His judgment that all crimson things are red is based on an exercise of the skill of imagining prototype samples of crimson and of red. According to Williamson, the role played by experience in Norman's knowledge that all crimson things are red is not purely enabling because Norman's reliable skill in imagining prototype samples of crimson and red goes beyond his understanding of the sentence 'all crimson things are red', and the acquisition of this additional skill involves experience:

"Norman's past experience did more than enable him to grasp proposition [R1]. It honed and calibrated his skills in applying the terms 'crimson' and 'red' to the point where he could carry out the imaginative exercise successfully." (Williamson 2013, p. 297)

In considering whether experience plays an evidential role in Norman's knowledge of R1, Williamson even goes so far as to concede that according to one possible interpretation of the example, experience indeed plays an evidential role. He writes:

"One interpretation of the example is that, although Norman's knowledge of [R1] does not depend on episodic memory, and he may even lack all episodic memory of any relevant particular colour experiences, he nevertheless retains from such experiences generic factual memories of what crimson things look like and of what red things look like, on which his knowledge of [R1] depends. [...] On this interpretation, Norman's colour experience plays an evidential role in his knowledge of [R1], thereby making that knowledge a posteriori." (Williamson 2013, p. 298)

Surprisingly, Williamson gives no reason for why this interpretation of the example is not the most plausible one.⁷ Instead, he merely claims that "we need not develop the example that way" (ibid.), and that the case might also be interpreted as one in which the only residue of Norman's experience active in his knowledge of R1 is a certain skill in recognizing and imagining (cf. ibid.). Williamson remarks that his choice of not interpreting the example as one of a posteriori knowledge is provisional, and that in section 5 of his paper, he will "reconsider, but reject, the idea that even supposed paradigms of a priori knowledge are really a posteriori" (cf. ibid.). However, in section 5 he merely argues that Quinean considerations regarding the interconnectedness of the epistemic statuses of individual beliefs do not provide good reasons for accepting the claim that all knowledge is a posteriori. This argument does not provide a convincing reason for not interpreting the particular case of Norman's knowledge of R1 as a case in which experience plays an evidential role. Moreover, as the above discussion of the role played by experience in Norman's knowledge of R2 demonstrates, the two possible

⁷ A very similar point is made by Carrie Jenkins in a short (and generally sympathetic) discussion of the earlier version of Williamson's argument in *The Philosophy of Philosophy* (cf. Jenkins 2008, p. 700).

interpretations that Williamson considers are not mutually exclusive. That is, analogously to what has been argued in the previous section, the supposed fact that the only residue of Norman's colour experience in his knowledge of R1 is a certain skill in recognizing and imagining does not guarantee that the role played by experience in his knowledge of R1 is less than evidential.

However, the fact that Williamson does not sufficiently argue for the claim that experience plays no evidential role in Norman's knowledge of R1 should not by itself be taken as a reason to reject it. In what follows, I will present an argument for the claim that experience plays no evidential role in Norman's knowledge of R1, and I will highlight the main shortcoming of this kind of argument.

Consider again the experience requirement in virtue of which, according to Williamson, the role played by experience in Norman's knowledge of R1 is more than purely enabling. The fulfilment of this requirement is exactly what distinguishes the case of Norman from that of Norbert, and it results in Norbert's being "less skilful than Norman in imagining a crimson sample." (Williamson 2013, p. 297) More exactly, the experience requirement in question concerns Norman's reliable skill of imagining not just some arbitrary sample of crimson, but a central prototype of such a sample, as well as some cognitive mechanism that enables Norman to reliably assess how much a colour shade may differ from the imagined one in order to still count as a shade of crimson. Analogously, in the case of Norman's knowledge of R2, the experience requirement in question concerns Norman's reliable skill of imagining the colour of recent volumes of *Who is Who*. But despite this superficial analogy, one might still suspect that the connection between the experience requirement in virtue of which, according to Williamson, the role played by experience in Norman's knowledge of R1 is more than purely enabling and Norman's understanding of the word 'crimson' is (in some way that would yet have to be specified) significantly closer than the connection between the experience requirement in virtue of which the role played by experience in Norman's knowledge of R2 is more than purely enabling and Norman's understanding of any of the expressions involved. That is, one might suspect that the connection between the relevant offline skills regarding the colour crimson and one's understanding of the corresponding term is in some way significantly closer than the connection between the reliable skill of imagining the colour of recent volumes of *Who is Who* and one's understanding of the terms 'recent' and 'volumes of *Who is Who*'. (One reason might be that my considerations in favour of the claim that experience plays an

evidential role in Norman's knowledge of R2 are not in any obvious way applicable to Norman's knowledge of R1.) And this suspicion might lead one to endorse the following argument:

- (P1) Some experience required for Norman's way of coming to know R2 plays an evidential role.
- (P2) Any experience required for Norman's way of coming to know R1 plays some non-evidential role (be it an enabling, or a non-enabling, but still less than evidential one) in his knowledge of R1.
- (C) Some of the experience required for Norman's way of coming to know R2, but none of the experience required for Norman's way of coming to know R1, plays an evidential role.

If this argument turned out to be sound, the considerations in favour of its soundness would provide a promising basis for defending the a priori/a posteriori distinction against Williamson's considerations. For given that experience plays an evidential role in Norman's knowledge of R2, but not in his knowledge of R1, this difference might appear epistemically significant enough in order to call Norman's knowledge of R2 a posteriori without risking that there will be no place left for the a priori. If Williamson is right in claiming that the role played by experience in Norman's knowledge of R1 is more than purely enabling, and if he is also right in claiming that the same holds for any knowledge whatsoever, then one might still use the term 'a priori' for all instances of knowledge in which experience does not play an evidential role – thereby using the terms 'a priori' and 'a posteriori' in order to mark a fundamental epistemological distinction. If, on the other hand, Williamson is right in claiming that the role played by experience in Norman's knowledge of R1 is more than purely enabling, but wrong in claiming that the same holds for all other supposed instances of a priori knowledge, then the theoretically most useful way of applying the terms 'a priori' and 'a posteriori' might be to reserve the former for those instances of knowledge in which experience plays a purely enabling role, while classifying Norman's knowledge of R1 as neither a priori nor a posteriori.

But before seriously engaging in this kind of pragmatic considerations, one should have a closer look at the above argument. Regardless of whether the premises of the argument can be convincingly defended, a closer look reveals that the argument is not

even valid. That is, even if P1 and P2 are true, C is still not established. In order to make the argument valid, the following additional premise is needed:

(P3) If some experience required for Norman's way of coming to know plays some non-evidential role (be it an enabling, or a non-enabling, but still less than evidential one) in his knowledge of R1, then it does not also play an evidential role in his knowledge of R1.

Adding P3 makes the argument valid. However, it is far from obvious that P3 is true. That is, it is far from obvious that the following claim is false.

(DR) Some of the experience that plays a non-evidential role in Norman's knowledge of R1 also plays an additional, evidential, role in his knowledge of R1.

Establishing DR (short for 'double-role') would not only refute P3, it would also, taken together with my above considerations regarding Norman's knowledge of R2, provide a strong case for Williamson's claim that there is no significant epistemological difference between Norman's knowledge of R1 and his knowledge of R2. In the following section, I will present a strategy for arguing in favour of DR.

1.4. The argument from analogy

As already announced at the beginning of the previous section, the argumentative strategy that I will now present consists in drawing an analogy between Williamson's version of the scenario of Norman's coming to know R1 and a modified version of it. The modified version has apparently not been considered by Williamson, but the argumentative strategy in question rests, at least to some extent, on Williamson's relatively minimal conception of understanding.

The modified version that I have in mind is one in which, in contrast to the original version, Norman does not learn the term 'crimson' by learning to distinguish crimson from non-crimson things on the basis of superficial colour perceptions. He has seen many crimson things under normal perceptual circumstances, and he has thereby had

the relevant superficial colour perceptions. But he has never been taught that these things are crimson. Instead, he learned the word ‘crimson’ from a scientist friend, Mary⁸, who told him all there is to know about the microscopic surface properties that all and only crimson things possess.⁹ Norman thereby learned to distinguish crimson from non-crimson things, not on the basis of superficial colour perception, but on the basis of observing the surfaces of the objects in question through a microscope. He has thereby also acquired good offline reasoning skills regarding the microscopic surface properties of crimson things. However, Mary has not told him about the microscopic surface properties that all and only red things possess. What Mary has told him has therefore not put him in a position to know, without having further experiences, that all crimson things are red. But even though Norman has not learned the word ‘crimson’ on the basis of superficial colour perceptions, he *has* learned the word ‘red’ in this way, and he has also acquired good offline reasoning skills regarding the superficial looks of red things. Obviously, though, his offline reasoning skills regarding the microscopic surface properties of crimson things and his offline reasoning skills regarding the superficial looks of red things do not yet put him in a position to know, without having further experiences, that all crimson things are red. In order to acquire this knowledge, he would (if he does not want to rely on testimonial evidence) need to have further experiences either regarding the microscopic surface properties of red things, or regarding the superficial looks of crimson things. He then finally comes to know R1, not through having further experiences regarding the microscopic surface properties of red things, but through having further experiences regarding the superficial looks of crimson things.

What role could this example play in an argument for DR? The first main step in such an argument could be the following claim:

⁸ This is of course an allusion to Frank Jackson’s famous thought experiment, presented in his 1982, p. 130, and in his 1986, p. 291. Williamson mentions Jackson’s thought experiment in a very short endnote (cf. Williamson 2013, p. 310, n. 6). But, analogously to his dismissal of the claim that experience plays an evidential role in Norman’s knowledge of R1, he does not offer any argument for why experience does not play an evidential role in Mary’s knowledge of what red things look like, apart from suggesting that examples like these might also be interpreted as ones in which the only residue of colour experience is the relevant skill of recognizing and imagining colours (cf. *ibid.*, p. 296).

⁹ Actually, mapping colours onto surface properties is not a straightforward matter, since similarity of colour appearance for relevantly similar subjects under relevantly similar visual circumstances does not guarantee similarity of surface properties (cf. Hardin 1988, pp. 5f.). However, this fact is of course still compatible with the claim that colour supervenes on surface structure. When I speak of surface properties that *all* and only things of a certain colour possess, these properties are supposed to be understood disjunctively. An attempt at mapping colours onto electromagnetic reflectance profiles, which, in turn, supervene on molecular surface structure, is offered in Churchland 2007.

(P1) In virtue of possessing the reliable skill of correctly applying the word ‘crimson’ on the basis of observing the surfaces of the objects in question through a microscope, Norman understands the word ‘crimson’.

That is, he understands the word before he comes to know what crimson things superficially look like. Williamson’s own remarks about the conditions for understanding can be taken as speaking in favour of this claim. In his 2006 paper “Conceptual Truth”, Williamson suggests that understanding a word requires no more than “causal interaction with the social practice of using that word [permitting] sufficiently fluent engagement in the practice.” (Williamson 2006, p. 38) Regarding what a sufficiently fluent engagement might consist in, he merely claims that it “can take many forms, which have no single core of agreement” (ibid.).¹⁰ In the present example, a causal connection to the social practice of using the word ‘crimson’ is given, since Mary came to use this word through interaction with people who learned the word in a perfectly normal way. But is Norman’s engagement in the practice of using the word sufficiently fluent in order to constitute an understanding of it? Note that, at the point of time in question, even though Norman only knows about the microscopic surface properties of crimson things, he already knows that crimson is a *colour*, and not, for example, a degree of smoothness. This kind of knowledge might already be taken as enabling a fluent engagement in the practice of using the word. Moreover, note that Norman is aware of the fact that he is not able to distinguish crimson from non-crimson things on the basis of superficial looks. A person who knows that ‘crimson’ is a colour term, but who applies the term, on the basis of superficial looks, to those things that superficially look green, might plausibly be regarded as not understanding the word ‘crimson’. Norman, however, is clearly better off in this regard. If he uses the term

¹⁰ Williamson reuses the same formulations quoted here in *The Philosophy of Philosophy*, p. 126, in a chapter that is, like the paper “Conceptual Truth”, primarily concerned with epistemological conceptions of analyticity. Williamson’s view on this topic is that there are no analytic truths in this sense, since he does not believe in any kind of direct understanding-assent link. His denial of the existence of these links relies on the claim that one might understand a sentence (or grasp a thought), and yet, due to deviant logical views or dispositions, not be disposed to assent to it. See Boghossian 2011 for objections against this latter claim, and see Williamson 2011 for a detailed response. These issues, however, are not relevant for present purposes, for even if Williamson is right in denying the existence of direct links between understanding and dispositions to assent, the requirement of possessing the right logical views or dispositions is not one that could tip the scales regarding whether Norman’s knowledge of R1 is a priori or a posteriori.

‘crimson’, then he does not usually apply it to the wrong things.¹¹ Instead, he is even able to explicitly teach normal speakers something about the essential characteristics of crimson things. In the light of these considerations, the claim that Norman understands the word ‘crimson’, even before getting to know what crimson things superficially look like, does not seem very far-fetched. It would of course take a lot more to conclusively establish it.

The next main step in an argument for DR could be the following claim:

(P2) If some experiences necessary for a certain way of coming to know a certain truth play neither the role of providing for the subject’s understanding of the respective sentence nor the role of providing for the subject’s offline reasoning skills corresponding to the subject’s way of understanding the respective sentence, then these experiences play an evidential role in the subject’s knowledge of that truth.

If one accepts P1, then one would also have to accept that the description in the antecedent of P2 is true of at least some of Norman’s experiences regarding the superficial looks of crimson things. For having these experiences is necessary for his way of coming to know R1, and before he has these experiences, he already understands, according to P1, the word ‘crimson’, as well as (so one might simply stipulate) the rest of the sentence, and he, according to the description of the scenario, already possesses good offline reasoning skills regarding the microscopic surface properties of crimson things and the superficial looks of red things. That is, if one accepts P1, one would also have to accept that at least some of Norman’s experiences regarding the superficial looks of crimson things neither play the role of providing for his grasp of R1 (or for his understanding of the corresponding sentence¹²) nor the role of providing for his offline reasoning skills corresponding to his way of grasping R1. Therefore, the following can be inferred from P1 and P2:

¹¹ For related reasons, P1 is, for instance, not in tension with the claim that “someone who sincerely denies that nothing can be red and green all over at once seem[s] to exhibit an inadequate understanding of at least one crucial term used in expressing that proposition“ (Audi 1998, p. 108). For if Norman had learned the term ‘red’ in the unusual way described above, while having learned the term ‘green’ in the usual way, he would, *ceteris paribus*, still not *deny* that nothing can be red and green all over at once.

¹² In later chapters (especially in chapter 3), I will be more specific regarding the objects of grasp or understanding that are relevant within the context of debates about apriority.

(C1) In the modified version of the scenario, some of Norman's experiences regarding the superficial looks of crimson things play an evidential role in his knowledge of R1.

In the spirit of Williamson's original argument, one might object to C1 by claiming that the only residue of Norman's experiences of the superficial looks of crimson things active in his knowledge of R1 is his skill of recognizing and imagining crimson things. However, this supposed fact would, by itself, not be a good reason for rejecting C1, because, as argued in section 1.2 above, reliably exercising the skill of imagining certain things may require experience to play an evidential role. Moreover, in the modified version of the scenario, Norman *already possesses* the skill of recognizing and imagining crimson things by the time he starts to have experiences regarding the superficial looks of crimson things. That is, he is already able to recognize crimson things by observing them through a microscope and to imagine crimson things by visually representing to himself their microphysical surface structure.

If one still wishes to insist on the claim that in the modified version of the scenario, Norman's experiences regarding the superficial looks of crimson things play no evidential role in his knowledge of R1, and if one also accepts P1, then one could be pressed to accept certain claims which defenders of the a priori will probably be quite unwilling to accept. Think, for example, of someone who comes to understand the word 'transparent' in a perfectly normal way, that is, by learning to distinguish transparent from non-transparent things on the basis of their superficial visual properties. Suppose further that this person comes to understand the term 'H₂O' in chemistry class, by learning to distinguish H₂O from other substances on the basis of looking through a microscope. She has, at that point, not been told that H₂O is water, and she has not explicitly been presented with H₂O except when looking through a microscope. She therefore does not know that H₂O is transparent. In order to acquire this knowledge without relying on testimony, she would need to have experiences regarding the superficial visual properties of H₂O (or regarding the microphysical properties of transparent things). Now, the following problem arises for those who deny C1 while accepting P1. Given that Norman understands the word 'crimson', and given that both Norman and the person in the example just described possess good offline reasoning skills corresponding to their ways of understanding the terms 'crimson' and 'H₂O', respectively, someone who denies C1, would, by analogy, apparently also commit

herself to the claim that, in the example just described, the person's experiences regarding the superficial visual properties of H₂O do not play an evidential role in her knowledge that H₂O is transparent. Those who deny C1, while accepting P1, would therefore apparently be committed to accepting the undesirable consequence that in the example just described, experience plays no evidential role.

This defense of C1 against someone who accepts P1 in effect amounts to a defense of P2. I am, of course, not claiming that the above considerations conclusively establish the truth of P2. (They will be dealt with in more detail in chapter 2.) As with P1, a lot more would be required for that. However, accepting P1 and P2 clearly commits one to accepting C1.

So far, the argument is merely concerned with the *modified* version of the scenario. As mentioned above, the purpose of presenting the argument is to present a strategy for establishing the truth of DR. But DR is concerned with the *original* version of the scenario. Accordingly, the final part of the argument consists in drawing an analogy between the two versions of the scenario. Before doing so, one should emphasize some of the similarities between the two versions. In both of them, Norman comes to know R1. And he comes to know it via the same method, namely, by representing to himself the superficial colour properties of crimson things and red things. Moreover, in both versions of the scenario, the reliable exercise of this offline skill partly stems from Norman's visual experiences regarding the superficial looks of crimson things. If the argument from P1 and P2 to C1 is sound, then these experiences play an evidential role in the modified version of the scenario. But given the similarities just mentioned, it might seem odd to accept C1 while denying that Norman's experiences regarding the superficial looks of crimson things also play an evidential role in the *original* version of the scenario. Drawing this analogy constitutes the final premise of the argument:

(P3) If C1 holds, then Norman's experiences regarding the superficial looks of crimson things also play an evidential role in his knowledge of R1 in the original version of the scenario.

Applying modus ponens to C1 and P3 yields the following conclusion regarding the original version of the scenario:

(C2) In the original version of the scenario, Norman's experiences regarding the superficial looks of crimson things play an evidential role in his knowledge of R1.

This completes my presentation of a possible strategy for establishing the claim that in the original version of the scenario, some experiences play the kind of double-role described in DR. Of course, DR does not directly follow from C2. But it may simply be granted to defenders of the apriority of R1 that Norman's experiences regarding the superficial looks of crimson things play an enabling role in his knowledge of R1; and the conjunction of this latter claim and C2 entails DR.

The argument to C2 will hitherto be referred to as 'the argument from analogy', and all uses of 'P1', 'P2', etc., will refer to the respective elements of this argument (rather than to elements of the argument considered in section 1.3).¹³

In the parlance of someone who endorses the argument from analogy, the argumentative strategy behind it might broadly be described as red-flagging the evidential role of some of Norman's experiences by means of *isolating* this role in a scenario in which this role and the other role that the same experiences originally play are played by different experiences.

As with P1 and P2, I do not claim to have established the truth of P3. A lot more would be required for that. However, my short considerations regarding these three premises already indicate that they are not far-fetched, and that the argument from analogy therefore deserves some closer examination, which will be provided in the following chapters.

1.5. Summary and outlook

In the foregoing sections, I have argued that Williamson's argument for the claim that there is no significant epistemological difference between Norman's knowledge of R1 and his knowledge of R2 is unsound; but I have presented an argumentative strategy for vindicating this claim.

¹³ In order to make later references to the elements of the argument from analogy more reader-friendly, I list the premises and conclusions of the argument at the end of the dissertation, in appendix A.

In particular, I have argued, in section 1.2, that Williamson's argument for the claim in question is unsound by showing that experience plays an evidential role in Norman's knowledge of R2. In section 1.3, I have considered an argument for the claim that experience does not play an evidential role in Norman's knowledge of R1. It has turned out that this argument relies on the hidden premise that if some experience plays a role in Norman's knowledge of R1 which is not evidential, then this experience does not also play a second, evidential, role in Norman's knowledge of R1. In section 1.4, I have presented an argument for the claim that some experience *does* play a second, evidential, role in Norman's knowledge of R1. If this argument, which I have called the 'argument from analogy', is sound, then it might be taken to vindicate Williamson's claim that there is no significant epistemological difference between Norman's knowledge of R1 and Norman's knowledge of R2, for it would then have been established that experience does not only play an evidential role in Norman's knowledge of R2, but also in Norman's knowledge of R1.

Furthermore, modified versions of the argument from analogy apply equally well to truths like the one expressed by the sentence 'Nothing is red (all over) and green (all over) at the same time', which might be taken as an even clearer example of an a priori truth than R1 (even though this is probably only due to the fact that R1 is less obviously *true* than the sentence just mentioned). The corresponding scenario would be one in which Norman is able to identify red objects on the basis of their superficial looks, but not on the basis of their microscopic surface properties, and in which he is able to identify green objects on the basis of their microscopic surface properties, but not on the basis of their superficial looks. Given this kind of recipe for reapplying the argument from analogy, I take it to be clear that the argument, if successful, would, together with the assumptions regarding apriority mentioned in the introduction, constitute a successful case for my central claim that colour inclusion sentences do not express a priori truths. Here I further assume that if Norman's knowledge of R1 in the original version of the scenario is not a priori, then there cannot be a priori knowledge of R1. (Otherwise, R1 might still be regarded as an a priori truth.) Even though this assumption may not seem entirely harmless, significant parts of my argumentation in favour of P3 (especially my considerations in section 4.2) will be generalizable to any other way of coming to know R1.

My argumentation for my central claim that colour inclusion sentences do not express a priori truths mainly consists in the argument from analogy. In the chapters to

come, I will therefore provide a detailed defense of the argument's soundness. That is, I will defend the argument against certain objections, and I will provide positive considerations in favour of it. I start, in the following chapter, by considering a relatively self-suggesting objection against the argument to C1 – an objection primarily directed against P2.

2. Examining the argument from analogy, part I: degrees of understanding

In this chapter, I will consider a possible strategy for blocking the first part of the argument from analogy on the basis of the notion of degrees of understanding. In section 2.1, I will describe how this strategy is supposed to work. In section 2.2, it will be pointed out that a certain classification of ways of understanding a word has to be endorsed in order for the strategy not to overgeneralize. In section 2.3, I will consider different ways of grounding the required classification. In section 2.4, I will argue that the considered ways of grounding the required classification, and thereby the whole strategy of blocking the argument from analogy on the basis of the notion of degrees of understanding, ultimately fails.

2.1. Blocking the argument to C1

The strategy for blocking the argument from analogy that I will consider in the present chapter concerns the first half of this argument, that is, the inference from P1 and P2 to C1:

- (P1) In virtue of possessing the reliable skill of correctly applying the word ‘crimson’ on the basis of observing the surfaces of the objects in question through a microscope, Norman understands the word ‘crimson’.
- (P2) If some experiences necessary for a certain way of coming to know a certain truth play neither the role of providing for the subject’s understanding of the respective sentence nor the role of providing for the subject’s offline reasoning skills corresponding to the subject’s way of understanding the respective sentence, then these experiences play an evidential role in the subject’s knowledge of that truth.
- (C1) In the modified version of the scenario, some of Norman’s experiences regarding the superficial looks of crimson things play an evidential role in his knowledge of R1.

As it has been stipulated regarding the modified version of the scenario, Norman's way of coming to know R1 partly consists in coming to know what crimson things superficially look like. Given P1, the kind of experience necessary for coming to know what crimson things look like does not play the role of providing for Norman's understanding of the word 'crimson'. And accordingly, it also does not play the role of providing for Norman's offline reasoning skills corresponding to the way in which he comes to understand the word 'crimson'. Therefore, the deductive step from P1 and P2 to C1 should be regarded as unassailable (with a caveat to be mentioned below). Consequently, the strategy for blocking the argument from analogy that I will consider in the present chapter concerns the truth of the two premises. In particular, it mainly concerns the truth of P2.

The strategy to be considered now is based on the claim that understanding comes in degrees. That is, someone might understand something better, or more fully, than someone else, even though this other person is also not in the state of having no understanding at all of the thing in question. This claim is intuitively very plausible when it comes, for example, to understanding states of affairs, since utterances of the form 'Person A understands the situation better than person B' are quite common and appear to be perfectly felicitous. And the same holds regarding the understanding of events, motives, opinions, etc. It seems natural to suppose that linguistic understanding is no exception here. And independently of what holds regarding other sorts of understanding, it would be quite implausible to deny that Norman's coming to know what crimson things superficially look like in the modified version of the scenario improves, or enriches, his understanding of the word 'crimson'. Moreover, the expressions 'understanding the word w' and 'being linguistically competent with respect to w' are very often used interchangeably in the philosophical literature about linguistic understanding; and it is hardly deniable that in the modified version of the scenario, Norman becomes linguistically more competent with respect to the word 'crimson' by coming to know what crimson things superficially look like.

How do these observations relate to P1 and P2? The answer depends on whether the degree of linguistic understanding on the part of a subject S necessary for truthfully ascribing to S understanding *simpliciter* of the word in question is context-dependent or context-invariant. I will not take a stance regarding whether the correct conception of understanding is a contextualist or an invariantist one. Instead, I will argue that the claim that linguistic understanding comes in degrees provides a *prima facie* promising

ground for blocking the argument to C1, independently of whether one adopts a contextualist or an invariantist conception of understanding. However, the strategy for blocking the argument to C1 to be considered here takes slightly different forms depending on which of the two conceptions is adopted. I will first develop the strategy in an invariantist framework. Later on, I will develop it in a contextualist framework.

So how does the claim that linguistic understanding comes in degrees relate to P1 and P2, given an invariantist conception of understanding? The first thing to note is that it does not provide a good reason for rejecting P1. For in an invariantist framework, the mere claim that linguistic understanding comes in degrees does not provide a good reason for denying that someone whose understanding of the word in question can still be improved or enriched can nevertheless possess understanding *simpliciter* of the word in question. Claiming that Norman does not understand the word ‘crimson’ in the modified version of the scenario (before coming to know what crimson things superficially look like) simply because his understanding of the word can still be improved would not be significantly more plausible than claiming that someone does not understand a certain language simply because she does not understand every single word of it.

But even though the claim that linguistic understanding comes in degrees does not provide a good ground for rejecting P1, it provides a *prima facie* promising ground for rejecting P2. But matters are not perfectly straightforward here, for the assumption that linguistic understanding comes in degrees might suggest an interpretation of P2 along the following lines:

(P2_D) If some experiences necessary for a certain way of coming to know a certain truth play neither the role of *improving* or *enriching* the subject’s understanding of the respective sentence nor the role of providing for the subject’s offline reasoning skills corresponding to the subject’s way of understanding the respective sentence, then these experiences play an evidential role in the subject’s knowledge of that truth.

Given this interpretation of P2, C1 cannot be derived from P1 and P2 (– this is the caveat announced above –), since none of the experiences necessary for Norman’s coming to know R1 in the modified version of the scenario satisfy the antecedent of

P2_D. However, P2_D is not a natural interpretation of P2. Rather, I take it that P2 would naturally be interpreted as follows:

(P2_S) If some experiences necessary for a certain way of coming to know a certain truth play neither the role of providing for the subject's understanding *simpliciter* of the respective sentence nor the role of providing for the subject's offline reasoning skills corresponding to the subject's way of understanding the respective sentence, then these experiences play an evidential role in the subject's knowledge of that truth.

Given this interpretation of P2, C1 can be derived from P1 and P2. For given P1, some of the experiences necessary for Norman's way of coming to know R1 – namely those experiences that provide for Norman's knowledge of what crimson things superficially look like – satisfy the antecedent of P2_S.

However, the claim that linguistic understanding comes in degrees provides a *prima facie* promising ground for rejecting P2_S. In particular, a conception of understanding as occurring in degrees might suggest the following claim regarding evidentiality of experience:

(NE) If some experiences necessary for a certain way of coming to know a certain truth improve or enrich the subject's linguistic understanding of one of the relevant expressions – while the subject already possesses linguistic understanding *simpliciter* of each of the relevant expressions as well as good offline reasoning skills corresponding to this understanding –, then these experiences play no evidential role in the subject's knowledge of that truth.

There is an obvious tension between NE and P2_S. More specifically, the antecedent of NE entails the antecedent of P2_S, and the consequent of NE entails the falsity of the consequent of P2_S. These two facts about NE and P2_S do not yet render the two claims incompatible. However, the falsity of P2_S can be derived from the conjunction of (NE) and the following supplementary claim:

(NE₊) The antecedent of (NE) is possible.

More precisely, the conjunction of NE and NE₊ entails the falsity of P2_S only if both NE and P2_S are read as *strict*, and not just material, conditionals. But it should be clear that this has been the intended reading throughout.

Given the plausible assumption that, in the modified version of the scenario, Norman's coming to know what crimson things superficially look like improves or enriches his understanding of the word 'crimson', it follows from P1 and NE that the experience involved in Norman's coming to know what crimson things superficially look like does *not* play an evidential role in his knowledge of R1 in the modified version of the scenario. Thus, endorsing NE is incompatible with endorsing the argument to C1.

The strategy of blocking the argument to C1 on the basis of the claim that understanding comes in degrees has hitherto merely been developed in an invariantist framework. I will now go on to develop it in a contextualist framework.

Contextualism regarding linguistic understanding might, for the moment, be defined as the view that whether an utterance of the form 'S understands expression e at time t' expresses a true or a false proposition depends on the context of utterance, even though 'S', 'e', and 't' are used to refer to the same person, expression (of a certain language with a certain meaning), and time, respectively. This general characterization of the view is still compatible with different versions of it.¹⁴ For present purposes, however, these differences may be neglected. In order to provide the ground for a promising way of blocking the argument to C1, the preferred version will merely have to be one according to which there are contexts in which the utterance 'Norman understands the word 'crimson' in the modified version of the scenario' expresses a true proposition, and according to which there are contexts in which this utterance expresses a *false* proposition. One may assume that all plausible versions of contextualism regarding linguistic understanding will provide this result.

¹⁴ According to one such version, the term 'understand' may express different relations in different contexts of utterance, and it may thereby contribute different intensions to the propositions expressed. According to an alternative version, the truth-value of ascriptions of understanding depends on the standards for understanding in the respective context, even though the respective utterances of the term 'understand' contribute the same intension to the propositions expressed. MacFarlane 2007 calls these two kinds of contextualism indexical and non-indexical contextualism, respectively. Regarding ascriptions of understanding, neither of the two versions of contextualism has so far been explicitly endorsed. (In fact, defenses of contextualism regarding linguistic understanding are hardly to be found in the literature. A rare example is Stanley 1999. However, in the light of Stanley's later endorsement of an invariantist account of knowledge (see especially his 2005), which is compatible, but not really in harmony with contextualism regarding linguistic understanding, it may be assumed that he stopped endorsing the latter view.)

Accordingly, the first step towards blocking the argument to C1 in a contextualist framework is to point out that P1, conceived of as a sentence of English, does not express a true proposition in all contexts of utterance; and this is the case even if in all contexts of utterance taken into consideration, the name ‘Norman’ is used to refer to the person described in the modified version of the scenario, and the sentence is used to describe a state of affairs occurring at a time when Norman is able to distinguish crimson from non-crimson things on the basis of observing the surfaces of the objects in question through a microscope. For the sake of brevity, I will simply rephrase this first step towards blocking the argument to C1 as the claim that P1 is not true in all contexts. Contexts in which P1 is false do not have to be particularly *demanding* ones. In some of them, it might, for instance, already be sufficient for Norman to be able to have some imperfect skill of distinguishing crimson from non-crimson things on the basis of their superficial looks. These contexts should not be regarded as more demanding than some of the contexts in which P1 is true.

From the claim that P1 is not true in all contexts, it follows that the argument to C1 can only be sound in a limited range of contexts. For the purpose of the contextualist’s strategy for blocking the argument, it suffices to point out that this limited range of contexts is one in which the requirements for linguistic understanding are not extremely high, since in all of the contexts in this limited range, Norman counts as understanding the word ‘crimson’ even though he is not able to distinguish crimson from non-crimson things on the basis of their superficial appearance. So if the argument to C1 is to be sound in any context, P2 has to be true in at least some of the contexts within the limited range just described.

Contextualists regarding linguistic understanding might indeed claim that P2 is true in some contexts. These contexts are such that in order to be able to truthfully utter sentences of the form ‘S understands expression e’, S’s understanding of e has to be such that it cannot be improved anymore. (Contextualists might, of course, disagree on whether such contexts can exist.) P2 might be regarded as true in these contexts, since NE₊ is false in them, so that arguing from NE and NE₊ to the negation of P2 will not be successful in these contexts. However, contextualists should plausibly claim that within the limited range of contexts in which P1 is true, NE₊ is true as well. Moreover, within the limited range of contexts in question, NE – even though not vacuously true – seems to be not less *prima facie* plausible than from the perspective of an invariantist framework.

In sum, then, the contextualist version of blocking the argument to C1 on the basis of the claim that understanding comes in degrees consists in claiming that there is no context in which both P1 and P2 are true, since in all contexts in which P1 is true, P2 might, like in an invariantist framework, be blocked on the basis of endorsing NE and NE₊.

One might point out that the distinction between contextualism and invariantism is not exhaustive. Even though contextualists and invariantists regarding linguistic understanding disagree on whether the truth-value of propositions expressed by utterances of the form ‘S understands expression e at time t’ – where the referents of ‘S’, ‘e’, and ‘t’ are fixed – can differ from one context of utterance to another, both contextualists and invariantists claim that it is an absolute matter whether an utterance in a given context expresses a true or a false proposition. That is, contextualists and invariantists agree (even though this is not part of the provisional definition of contextualism given above) that whether an utterance expresses a true or a false proposition at most depends on the context of utterance, but not on the context of *assessment* of the utterance. Relativists deny this latter claim.¹⁵ Relativist accounts of knowledge ascriptions have been defended in Richard 2004 and MacFarlane 2005, and such accounts might equally well be proposed regarding ascriptions of understanding. However, it will not be necessary to describe in any detail how a relativist regarding linguistic understanding might block the argument to C1, since it is relatively obvious that the relativist strategy will be analogous to the contextualist strategy described above. That is, in analogy to the strategy proposed for the contextualist, the relativist might claim that there is no context of assessment in which both P1 and P2 are true (or more explicitly: that there are no assessment standards according to which both utterances of P1 and of P2 express true propositions), since in all contexts of assessment in which P1 is true, P2 might, once again, be blocked on the basis of endorsing NE and NE₊.

However, the strategy for blocking the argument to C1 presented in this section – namely, the strategy of endorsing NE and NE₊ – faces a serious problem. In the following section, I will describe what this problem consists in, and how it might be avoided.

¹⁵ Given a wide enough conception of contextualism, one might also characterize relativism as a radical form of contextualism. For present purposes, nothing hinges on that.

2.2. The danger of overgeneralization

As described in the previous section, endorsing NE and NE₊ provides a strategy for arguing that the experience involved in Norman's coming to know (in the modified version of the scenario) what crimson things superficially look like plays no evidential role in Norman's knowledge of R1. This strategy, however, faces the problem of overgeneralizing to kinds of experience that even the most fervent defender of a priori knowledge would probably be unwilling to characterize as playing a non-evidential role.

In order to see this, consider a scenario in which Norman – as in the original version construed by Williamson – is able to reliably distinguish crimson from non-crimson things on the basis of their superficial looks, and in which he possesses good corresponding offline reasoning skills. Suppose further that Mary then instructs Norman how to find out about the microscopic surface structure of crimson things, and Norman, following her instructions, comes to know about the microscopic surface structure of crimson things on the basis of looking through microscopes. More precisely, Norman acquires propositional knowledge of the form 'All and only crimson things possess microscopic surface property C' (where C will probably be a disjunctive property – see fn. 9 above). Even the most ardent defender of a priori knowledge will probably characterize Norman's knowledge of the microscopic surface properties of crimson things as a posteriori, and she will, accordingly, characterize the experiences involved in Norman's use of microscopes as playing an evidential role in Norman's knowledge. However, endorsing NE and NE₊ is in tension with such a characterization.

Admittedly, endorsing NE and NE₊ does not by itself commit one to characterizing the role played by the experience involved in Norman's use of microscopes as non-evidential. For one might claim that, even though Norman's understanding of the word 'crimson' could still be improved (e.g. via an improvement of Norman's skills of discriminating crimson from non-crimson things on the basis of their superficial looks), Norman's newly acquired knowledge about the microscopic surface structure of crimson things does not at all improve or enrich his understanding of the word 'crimson'. That is, one might claim that this kind of scientific knowledge is irrelevant regarding how well one understands the word. However, this claim is not plausible. For in the version of the scenario presented in section 1.4, in which Norman has no idea what crimson things superficially look like, and in which he is merely able to distinguish crimson from non-crimson things on the basis of their microscopic surface

properties, it is hard to deny that this scientific knowledge provides him with at least *some* understanding of the word ‘crimson’ – at least within a strategy, like the one considered here, which is based on a conception of understanding as coming in degrees. Moreover, it is hard to deny (within such a strategy) that this knowledge provides him with some understanding of the word ‘crimson’ that goes beyond Norman’s knowledge that ‘crimson’ is a colour term. But if this is so, then it should analogously be conceded, regarding the scenario described above, that Norman’s newly acquired knowledge (or probably already his justified belief, regardless of whether it amounts to knowledge)¹⁶ about the microscopic surface structure of crimson things improves or enriches Norman’s understanding of the word ‘crimson’ at least a bit, even if he already possesses understanding simpliciter in virtue of being able to distinguish crimson from non-crimson things on the basis of their superficial looks. One might, of course, still deny that understanding comes in degrees at all, and thereby deny that understanding simpliciter can be improved. But this claim is incompatible with NE₊. Therefore, denying that the experience involved in Norman’s use of microscopes improves his understanding of the word ‘crimson’ in the version of the scenario described above is no promising option for someone who tries to block the argument to C1 on the basis of endorsing NE and NE₊. But endorsing NE and NE₊, while conceding that the experience involved in Norman’s use of microscopes improves his understanding of the word ‘crimson’, commits one to the claim that this experience plays no evidential role in Norman’s knowledge that all and only crimson things possess microscopic surface property C – which, in turn, seems to leave no room for characterizing this knowledge as a posteriori. As already mentioned above, even the most fervent defender of a priori knowledge will probably be unwilling to accept this consequence. Defenders of the strategy might of course still flat out deny that, even though linguistic understanding

¹⁶ Pettit 2002 argues that linguistic understanding does not require knowledge of meaning, since – as he claims – only the latter, but not the former, fails in Gettier-style cases. Apparently independently of Pettit’s paper, Kvanvig 2003, pp. 197ff., presents the same kind of argument for a similar conclusion. One might transform these arguments into an argument to the conclusion that what is relevant regarding the degree of Norman’s understanding of the word ‘crimson’ is not the fact that he *knows* about the microscopic surface structure of crimson things, but only his corresponding justified belief, regardless of whether this belief amounts to knowledge. See Longworth 2008, pp. 56f., for objections against Pettit’s version of the argument, and see Grimm 2006 for objections against Kvanvig’s version of the argument. Pettit also argues that not even justification or belief, but only a certain kind of true seeming, is necessary for understanding the word in question. See Gross 2005 for objections against Pettit’s argument for this latter claim. For present purposes, these debates are irrelevant, since the kind of experience involved in Norman’s coming to know about the microscopic surface properties of crimson things – namely, the experiences provided by the use of microscopes – is also involved in Norman’s acquisition of the corresponding (justified) true (belief or) seeming.

comes in degrees, the experience involved in Norman's use of microscopes improves or enriches his understanding of the word 'crimson'. But I take it that, as long as such a denial is not supported by further argument, the strategy should be regarded as unconvincing.

Obviously, endorsing NE and NE₊ should (without further support for the denial just mentioned) be regarded as involving analogous commitments concerning paradigm examples of natural kind terms. The general pattern of revealing the kind of unwelcome consequence just described is the following. Whenever a kind *x* has a certain essential property *p*, such that it is necessarily the case that all and only things of kind *x* have property *p*, it is hard to deny that knowledge that all and only things of kind *x* have property *p* is relevant regarding how well one understands the term '*x*'. That is, as long as one accepts that linguistic understanding comes in degrees, it seems quite implausible to deny that coming to know that all and only things of kind *x* have property *p* increases one's degree of understanding of '*x*', even if one already possesses understanding simpliciter of the word. Examples of this kind of knowledge would be knowledge that all and only water has the chemical structure H₂O and knowledge that all and only gold is the element with atomic number 79 – to take the most paradigmatic examples of natural kinds. Endorsing NE and NE₊, while accepting that the respective knowledge increases one's understanding of the corresponding natural kind term, commits one to the claim that the kind of experience involved in coming to know about the chemical structure of water, or the atomic number of gold, plays no evidential role in one's knowledge that all and only water has the chemical structure H₂O, or that all and only gold is the element with atomic number 79, respectively. Analogously to the case of colour terms, this characterization of the roles of experience involved seems to leave no room for characterizing the cases of knowledge just mentioned as a posteriori.

Endorsing NE and NE₊ might even involve similarly unwelcome commitments regarding knowledge of contingent properties, if one accepts the respective claim that knowledge of the contingent fact that all and only things of kind *x* possess the property in question increases one's degree of understanding of the word '*x*'. However, I take it that the above considerations regarding essential properties already suffice for demonstrating that the strategy of blocking the argument to C1 on the basis of endorsing NE and NE₊ (without further argument for the claim that, even though understanding comes in degrees, the experience involved in Norman's use of microscopes does not improve or enrich his understanding of the word 'crimson') is unacceptable.

In order to save the strategy from this kind of objection, one would have to restrict NE in a certain way. The required kind of restriction would be based on a distinction between what may be called ‘primary’ and ‘secondary’ ways of understanding a term. The former would, by definition, have a privileged status regarding the question which of the experiences described in the antecedent of NE play an evidential role in the respective instance of knowledge. More exactly, those experiences required for a certain way of coming to know a certain truth that merely improve or enrich the subject’s *secondary* ways of understanding one of the relevant expressions – while the subject already possesses linguistic understanding simpliciter of each of the relevant expressions as well as good offline reasoning skills corresponding to this understanding – do, by definition, play an evidential role in the subject’s knowledge of that truth. This stipulation restricts the scope of NE in the following way:

(NE_{PR}) If some experiences necessary for a certain way of coming to know a certain truth merely improve or enrich the subject’s *primary* ways of understanding of one of the relevant expressions – while the subject already possesses linguistic understanding simpliciter of each of the relevant expressions as well as good offline reasoning skills corresponding to this understanding –, then these experiences play no evidential role in the subject’s knowledge of that truth.

Defenders of this modified version of the strategy also have to make certain claims concerning which ways of understanding are primary and which are secondary. In particular, they have to claim that understanding the word ‘crimson’ in virtue of being able to distinguish crimson from non-crimson things on the basis of their superficial looks is a *primary* way of understanding the word. Otherwise, NE_{PR} would be of no help for blocking the argument to C1. Defenders of the modified version of the strategy also have to claim that understanding the word ‘crimson’ in virtue of being able to distinguish crimson from non-crimson things on the basis of their microscopic surface properties, and understanding terms like ‘water’ and ‘gold’ in virtue of being able to distinguish portions of the respective stuff from portions of other stuff on the basis of their chemical composition, are *secondary* ways of understanding. Otherwise, the strategy would still overgeneralize in the way described above.

Without any further explanation of why some particular ways of understanding count as primary and others as secondary, this proposed solution to the problem of overgeneralization would appear quite ad hoc. In the following section, I will therefore describe what such an explanation might most plausibly look like.

2.3. Grounding the distinction between primary and secondary ways of understanding

As mentioned above, attempts to solve the problem of overgeneralization via the distinction between primary and secondary ways of understanding will be ad hoc as long as no explanation has been provided of why understanding the word ‘crimson’ in virtue of being able to distinguish crimson from non-crimson things on the basis of their superficial looks is a *primary* way of understanding the word, and of why understanding words like ‘crimson’, ‘water’, and ‘gold’ on the basis of microphysical or chemical properties are *secondary* ways of understanding. It seems quite plausible to assume that any promising attempt at providing such an explanation will be one according to which the fact that certain ways of understanding a word are primary is somehow reflected in the *content* of the respective word. Those who wish to deny this connection would – given how primary and secondary ways of understanding have been defined above – be committed to the claim that facts about whether NE is true of a certain way of understanding a given term are not reflected in the contents of the respective terms. I take it that a surprising story would have to be told in order to make a convincing case for this claim.

There are various options for grounding the above-mentioned classification of ways of understanding in the contents of the respective terms. One option is to claim that colour terms like ‘crimson’ have *intensions* – conceived of in a certain way to be described below –, and that the primary ways of understanding one of these terms are those that consist in grasping the intensions of the terms.

Intensions are nowadays most commonly defined as functions from possible worlds to extensions (or more explicitly: as functions from ways the world might have been to the things referred to by the respective term when the term is used to describe the ways the world might have been). It should, however, already be pointed out that the attempt at grounding the distinction between primary and secondary ways of understanding to

be considered here is one whose prospects are rather dim unless the relevant conception of intension is one according to which intensions are *not* merely functions to extensions, but rather some kind of cognitively individuated, descriptive entity. (I will provide an explanation of this at the end of the present section.) Following the principle of charity, I will therefore, in my description of the attempt at grounding the distinction between primary and secondary ways of understanding to be considered, employ a relatively rich notion of intension according to which intensions are cognitively individuated, descriptive entities, which may be described as containing (phenomenal) information.

Those attempting to ground the distinction between primary and secondary ways of understanding on the basis of the notion of intension would have to claim that the superficial looks of crimson things are somehow involved in the intension of the term 'crimson'. Understanding the word 'crimson' in virtue of being able to distinguish crimson from non-crimson things on the basis of their superficial looks could then be classified as a primary way of understanding the word, since it could be described as grasping (in some wide sense) the intension of the word. Accordingly, understanding the word 'crimson' in virtue of being able to distinguish crimson from non-crimson things on the basis of their microscopic surface properties could be classified as a secondary way of understanding the word, since it does not provide a grasp of the intension of the word.

As mentioned above, in order to avoid overgeneralization, defenders of the strategy of employing NE_{PR} also have to claim that ways of understanding supposedly clear examples of natural kind terms, like 'water' and 'gold', on the basis of the respective "hidden" properties are secondary ways of understanding these terms. One way of explaining this classification is to claim that these terms do not have intensions at all (given the rich notion of intension mentioned above), and that, accordingly, *any* way of understanding one of these terms is secondary. This claim, however, together with the above-mentioned claims regarding colour terms, would commit one to the claim that there is a fundamental semantic asymmetry between colour terms like 'crimson' and terms like 'water' and 'gold' (an asymmetry that has, of course, nothing to do with the superficial fact that the former is an adjective and that the latter are nouns). I take it that some argument would be required to justify the claim that there is such an asymmetry.

Another way of explaining the classification of ways of understanding terms like 'water' and 'gold' on the basis of the respective "hidden" properties as secondary is to claim that these terms do have intensions, but that the ways of understanding just

mentioned do not provide a grasp of these intensions. This option is quite problematic. For any promising candidates of informative intensions of terms like ‘water’ and ‘gold’ are incompatible with the (among philosophers) widely accepted view that theoretical identifications like ‘water = H₂O’ express necessary truths. That is, for any candidate for an intension of terms like ‘water’ and ‘gold’ which does not include the respective hidden properties, there will be possible worlds with respect to which the intension picks out something as water or gold that is not H₂O, or that does not have atomic number 79, respectively – namely those substances that have the superficial qualities included in the term’s intension, but a different hidden structure (see Putnam 1975, pp. 223ff., for the classic example).

Whether claiming that colour terms have the relevant kind of intension (that is, most plausibly intensions picking out those things that look a particular way to normal human beings under normal circumstances) involves analogous commitments regarding the modal status of the kind of theoretical identifications that Mary has taught Norman in the example in section 1.4 depends on what “normal circumstances” are supposed to include. If they do not only include lighting conditions, but also all relevant natural laws, then it might well be that the proposed kind of intension is compatible with the claim that the kind of theoretical identifications that Mary has taught Norman are necessarily true. But regardless of this rather complicated matter, defenders of NE_{PR} who wish to explain the required classification of particular ways of understanding as primary or secondary in terms of intensions are, as stated above, either committed to claiming that there is a fundamental semantic asymmetry between colour terms and terms like ‘water’ and ‘gold’, or they are committed to rejecting the widely accepted view that theoretical identifications like ‘water = H₂O’ express necessary truths.

I do not wish to claim that neither of these two commitments can be coped with. Instead, I will now turn to a more modest way of explaining the required classification of particular ways of understanding as primary or secondary in terms of intensions – a way that involves neither of the two commitments just mentioned. Later on, I will show that both ways of explaining the required classification – the one already described and the more modest one to be described below – ultimately fail for the same reason.

The more modest way in question consists in endorsing a view called two-dimensionalism. One of the central aims of this view is to reconcile the supposed Kripkean insight that there are identity statements which are both a posteriori and metaphysically necessary with the intuition that the contents of the expressions which

figure on the left and on the right side of these statements differ. Frank Jackson – one of the two main proponents of two-dimensionalism – formulates this latter intuition as follows:

“What we say about what the world is like using the sentence ‘There is water’ is different from what we say about what the world is like using the sentence ‘There is H₂O’. Otherwise it would not have been a *discovery* that water is H₂O. But the worlds at which the sentences are true are one and the same.” (Jackson 2004, p. 262)

So one of the central aims is to make sense of the idea that even though, for example, water is necessarily H₂O, the content of the word ‘water’ differs from that of the word ‘H₂O’. As mentioned above, this task involves the difficulty that as soon as one assigns to the term ‘water’ an intension including some information about the superficial qualities of water, one can no longer consistently hold that water is necessarily H₂O, since there will be some possible worlds with respect to which this kind of intension picks out a stuff which is not H₂O.

As the name of the theory already suggests, the proposed solution consists in distinguishing between two dimensions of content. That is, according to two-dimensionalism, each referring expression has two intensions of a fundamentally different kind. In his 1996 book *The Conscious Mind*, pp. 60ff., David Chalmers – the other of the two main proponents of the view – uses the terms ‘primary intension’ and ‘secondary intension’ for this distinction. I will adopt this terminology here.¹⁷ The roles for these two aspects of content are supposed to be as follows. While the secondary intension is that aspect of the content which accounts for the *necessity* of necessary a posteriori truths such as ‘water = H₂O’, the primary intension accounts for the aposteriority of those truths. Moreover, the primary intension is supposed to account for the alleged fact that there are certain a priori truths about water, while the secondary intension is supposed to account for the fact that these allegedly a priori truths are only contingently true. Chalmers 1996 suggests that what can be known a priori about water is roughly that it is “the dominant clear, drinkable liquid in the oceans and lakes” (p. 57), while conceding that such a formulation of the primary intension will probably not be accurate for those cases in which several different liquids occur to an equal extent (*ibid.*, p. 57f.). Jackson 1998 claims that what can be known a priori about water is that

¹⁷ Jackson 1998 uses the terms ‘A-intension’ and ‘C-intension’, but he also explicitly states that these terms are supposed to express the same distinction that Chalmers has in mind (cf. Jackson 2004, p. 261); Chalmers 2006 also employs the terms ‘first-dimensional intension’ and ‘second-dimensional intension’, and ‘1-intension’ and ‘2-intension’, to make the very same distinction.

it is “the watery stuff of our acquaintance” (p. 52), remaining unspecific about the properties in virtue of which a stuff counts as watery. In contrast, the secondary intension of ‘water’ includes a posteriori information about the kind of chemical substance that the primary intension picks out in the actual world.

In order to have a full grasp of the fundamental difference between primary and secondary intensions, it is helpful to have a look at the procedures necessary for revealing the primary and secondary intension of a term. According to two-dimensionalism, a term has its particular secondary intension partly in virtue of its particular primary intension, and not vice versa. Accordingly, the procedure of revealing the primary intension of a term does not in principle depend on any information about the term’s secondary intension. Rather, the procedure is supposed to be a purely a priori exercise. More precisely, it consists in considering an epistemic possibility (that is, a state of the world compatible with all one can know a priori) *as actual*. Chalmers & Jackson 2001 put this as follows:

“Let us say that a hypothesis is epistemically possible (in the broad sense) when it is not ruled out a priori. Let us say that an *epistemically possible* hypothesis characterizing the total state of the world corresponds to an *epistemic possibility*: intuitively, a specific way the actual world might turn out to be, for all one can know a priori.^{FN} Then sufficient information about an epistemic possibility enables a subject to know what a concept’s extension will be, under the hypothesis that the epistemic possibility in question is actual.” (Chalmers & Jackson 2001, p. 324)

In the case of ‘water’, it is supposedly compatible with everything competent users of this term can know a priori that the dominant watery stuff in the environment is actually not H₂O, but a substance with a different chemical structure. Chalmers & Jackson claim that if this possibility is actual, then water is nonetheless the dominant watery stuff in the environment:

“[I]n the case of water, given appropriate information about the distribution, behavior, and appearance of clusters of XYZ molecules (information analogous to the information we have about H₂O in the actual world), a subject is in a position to conclude that if the information is correct, then water is XYZ.” (Chalmers & Jackson 2001, p. 324)

The method of revealing the primary intension of ‘water’ can thus be formulated as follows: Consider various epistemic possibilities and observe, via a priori reasoning, those properties that water has in all those possible states of the world, under the assumption that those possibilities are actualities! The primary intension of ‘water’ includes exactly those properties that water has in all epistemically possible states of the world. (Employing the more standard definition of intension already mentioned above,

according to which intensions are mere functions from possible worlds to extensions, the primary intension of a term would be defined as the function from epistemic possibilities to the extensions of the term in these possible states of the world, under the assumption that those possibilities are actualities.) Difficulties for this kind of definition arise with respect to those epistemic possibilities in which the term has been assigned a clearly different meaning; but these details are not important here.

As already mentioned, the secondary intension of a term is supposed to be revealed on the basis of the term's primary intension. The procedure consists in observing (via a posteriori research) which kind of stuff (in the case of natural kind terms) is actually picked out by the primary intension of the term in question. Of course, in order to competently reveal a term's secondary intension, it is not necessary to have actually followed the above-mentioned procedure of revealing the term's primary intension. Rather, this latter procedure is merely required for having an *explicit* grasp of the term's primary intension. The procedure of revealing a term's secondary intension can also be reliably exercised solely on the basis of an implicit grasp of the primary intension. (Chalmers 2002, p. 43, speaks of intensions in general as "tacit criteria for identifying the extension of [an] expression".) In the case of natural kind terms like 'water' and 'gold', once we have found out about the kind of stuff actually picked out by the primary intension of the term, we are, in accordance with the Kripkean picture of natural kind terms, in a position to know that the term refers to this kind of stuff with respect to all possible worlds in which this stuff exists, at least as long as these possible worlds are not considered as actual, but merely as possible.¹⁸

Corresponding to the notions of primary and secondary intension, two-dimensionalists also distinguish between the primary and secondary *extension* of a term – the former being the set of things picked out by the term's primary intension, and the latter being the set of things picked out by the term's secondary intension. While the two-dimensionalists' use of the term 'primary intension' captures the traditional philosophical notion of intension, the traditional philosophical notion of extension is captured by their use of the term 'secondary extension'. Even though the supposed fact

¹⁸ The distinction between primary and secondary intensions basically rests on the same idea as several other distinctions which have appeared earlier in the literature, like those between diagonal propositions and propositions expressed (Stalnaker 1978), deep and superficial necessity (Evans 1979), and character and content (Kaplan 1989), even though the corresponding approaches are not supposed to be applied to Kripkean examples of a posteriori necessity. Davies and Humberstone 1980 provide a formal analysis of Evans' distinction between deep and superficial necessity, involving a modal operator they call 'fixed actually', and they are the first who explicitly employ the notion of considering a world as actual. For a detailed description of these and other predecessor-notions, see Chalmers 2006, pp. 112ff.

that the primary intension of the term ‘water’ includes properties like transparency, drinkability, a certain distribution, etc. is not compatible with the term’s primary extension being H₂O (for with respect to some possible worlds, such an intension picks out a different kind of stuff), it is still compatible with the term’s *secondary* extension being H₂O. The two-dimensionalists’ distinction between different aspects of content therefore allows them to reconcile the supposed apriority of certain descriptivist claims regarding terms like ‘water’ with Kripkean claims about the necessity of certain a posteriori truths regarding these terms. Or as Jackson & Chalmers put it:

“[N]othing here conflicts with the conclusions of Kripke (1972) [see Kripke 1980 in the bibliography] about names and natural kind terms. On Kripke’s picture, it is an a posteriori necessity that water is H₂O. But this is entirely compatible with there being an a priori conditional from certain (false) statements about the distribution, behavior, and role of XYZ to ‘water is XYZ.’” (Jackson & Chalmers 2001, pp. 325f.)

It is now obvious how two-dimensionalists can ground the distinction between primary and secondary ways of understanding in the content of the respective term. Primary ways of understanding a word will consist in an implicit or explicit grasp of the word’s primary intension. This yields the following version of NE:

(NE_{PR-INT}) If some experiences necessary for a certain way of coming to know a certain truth merely improve or enrich the subject’s grasp of the primary intension of one of the relevant expressions – while the subject already possesses linguistic understanding simpliciter of each of the relevant expressions as well as good offline reasoning skills corresponding to this understanding –, then these experiences play no evidential role in the subject’s knowledge of that truth.

As Nimtz 2004 points out, it is quite doubtful whether the above-mentioned procedure of revealing the primary intension of terms like ‘water’ will yield a set of properties that include superficial characteristics like transparency, drinkability, etc., since – so he argues – for any of those superficial properties, there will be a possible world in which some stuff has all the superficial properties associated with ‘water’, except the one property in question:

“[I]magine 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*.” (Nimtz 2004, p. 138)

Nimtz concludes that if two-dimensionalists wish to claim that primary intensions of terms like ‘water’ include the kind of properties just mentioned, they would have to resort to a disjunctive account of those primary intensions, which would commit them to accepting a striking asymmetry between natural kind terms and supposed nominal kind terms like ‘knowledge’, for “[i]n the latter case, we have no reason to believe that our tacit criteria for applying the term are disjunctive” (p. 139).

However, in order to block the argument to C1 on the basis of NE_{PR-INT} , one does not have to claim that the primary intensions of terms like ‘water’ and ‘gold’ contain any of the properties just mentioned. In order to avoid the kind of overgeneralization described in the previous section, it is merely required that a posteriori information regarding “hidden” properties is not included in the primary intension. What proponents of the strategy of blocking the argument to C1 on the basis of NE_{PR-INT} *do* have to claim, though, is that the primary intension of colour terms includes information about the superficial looks of things of the respective colour (under normal circumstances). It might, however, even after considering Nimtz’ objection, indeed seem plausible that the procedure of revealing the primary intension of colour terms will yield such a result. And in order to avoid being committed to accepting a fundamental asymmetry of the kind described above between terms like ‘water’ and ‘gold’, on the one side, and colour terms, on the other, one merely has to claim that the former have *some* kind of primary intension, which might merely include the property of being (or having been) called ‘water’ or ‘gold’, respectively, by a certain group of speakers. (Such sparse candidates for the primary intensions of words like ‘water’ and ‘gold’ might not be what two-dimensionalists originally hoped for; but this should not be relevant for present purposes.)

So even though two-dimensionalism provides a thinner way of grounding the distinction between primary and secondary ways of understanding than those ways of grounding that do not involve a differentiation between different kinds of content (since the former only provides a way of grounding the distinction in one *aspect* of the content), it avoids the above-mentioned problems that those who employ an undifferentiated notion of intension are faced with.

Having introduced the notions of primary and secondary intension, and having described the possible two-dimensionalist way of grounding the distinction between

primary and secondary ways of understanding, I can now – as already announced above – explain why the notion of intension employed by proponents of the strategy in question should not be the most standard one, according to which intensions are mere functions from possible worlds to extensions, but rather one according to which they are some kind of cognitively individuated, descriptive entity. The reason is the following.

In order to ground a distinction between primary and secondary ways of understanding terms like ‘water’ and ‘gold’, the standard conception of intensions might indeed be suitable, since the usual kinds of candidates for the primary intensions of these terms certainly map different extensions to possible worlds than the usual kinds of candidates for the *secondary* intensions of these terms. That is, it would have been possible for the watery stuff around here, or for the metallic, golden stuff around here (to use very rough candidates for the primary intensions of these terms), not to be H₂O or the element with atomic number 79, respectively. However, matters are different regarding the term ‘crimson’. A self-suggesting candidate for a primary way of understanding this term would be a way of understanding it by associating the term with a certain phenomenal appearance for perceivers of a certain kind under visual circumstances of a certain kind. And the secondary way of understanding the term would most obviously consist in associating it with a certain disjunctive microphysical surface property, or with a certain non-disjunctive light-emission property.

In contrast to the cases of ‘water’ and ‘gold’, one may plausibly hold that the functions corresponding to these two ways of understanding the term ‘crimson’ map exactly the same extensions to possible worlds. That is, it seems that it would not have been possible for things that phenomenally appear crimson to the relevant kind of perceiver under the relevant kind of visual circumstance not to have the disjunctive microphysical surface property or non-disjunctive light-emission property in question. And conversely, it seems that it would not have been possible for things with the disjunctive microphysical surface property or non-disjunctive light-emission property in question not to phenomenally appear crimson to the relevant kind of perceiver under the relevant kind of visual circumstance. Of course, these non-possibility judgments are only plausible given a specific enough conception of the relevant kind of perceiver and circumstance. But given such a specific enough conception, the functions corresponding to the self-suggesting candidates for the two ways of understanding the term ‘crimson’ are simply identical. Thus, in order to ground the distinction between a primary and a secondary way of understanding the term ‘crimson’, two-dimensionalists apparently

have to endorse a richer conception of intensions, according to which intensions are not mere functions to extensions, but rather some kind of cognitively individuated entity – thereby guaranteeing the non-identity of the primary and secondary intension of the term ‘crimson’.

Endorsing the kind of conception of intensions needed for the two-dimensionalist’s attempt at grounding the distinction between a primary and a secondary way of understanding the term ‘crimson’ might well be regarded as an unwelcome commitment. (Even Chalmers himself explicitly endorses the more standard, sparse conception of intensions – see, for instance, his 2002, p. 145, and 2004, p. 156). However, even setting this worry aside, the strategy fails for another reason, which I now turn to.

2.4. Why the strategy ultimately fails

As described in the previous section, explaining the required classification of a certain way of understanding the term ‘crimson’ as primary on the basis of the notion of primary intension is more promising than explaining it on the basis of an undifferentiated notion of intension. As I will argue in this section, however, both approaches fail, and they do so for the same reason.

As mentioned earlier, those who try to block the argument to C1 on the basis of some restricted version of NE would – if they do not wish to deny the connection between roles of experience and the content of the relevant expressions – have to claim that the superficial looks of crimson things (under normal circumstances) are somehow included in the (primary) intension of the word ‘crimson’. The reason why such a claim, as long as it remains unsupported by further arguments, is an illegitimate move within a strategy for blocking the argument to C1 becomes apparent as soon as one looks again at the scenario that triggered the argument from analogy. In this scenario, Norman is taught about the microscopic surface properties of crimson things, and he also acquires good offline reasoning skills regarding these properties. According to P1, he thereby comes to understand the word ‘crimson’, even though he still has no idea what crimson things superficially look like. Now consider the procedure which – as defenders of the strategy of blocking the argument to C1 on the basis of some restricted version of NE would have to claim – reveals that the superficial looks of crimson things (under normal

circumstances) are included in the (primary) intension of the word 'crimson'. This procedure is supposed to work as follows. Consider various possible worlds (as actual) and observe, via a priori reasoning, which properties of crimson things are the resilient ones! The claim that this procedure reveals that the (primary) intension of the word 'crimson' includes certain superficial looks (under normal circumstances) involves the claim that the superficial looks of crimson things (under normal circumstances) can be determined via a priori reasoning. Regarding the superficial properties supposedly included in the primary intension of the word 'water', Jackson formulates the alleged a priori relation as follows: "This is what we grasp when we come to understand the word." (Jackson 1998, p. 49) According to P1, however, Norman already understands the word 'crimson' before learning about the superficial looks of crimson things. And he is clearly not in a position to grasp the (primary) intension of the word 'crimson', if it is to include information about the superficial looks of crimson things.

Hence, the above-considered ways of explaining the required classification of certain ways of understanding as primary rely on the claim that P1 is false. Since these explanations are supposed to figure in a strategy for blocking the argument to C1, and since the reliance on the claim that P1 is false is not supported by any argument, they are ultimately circular. This is, of course, not to say that the strategy considered here cannot be supplied with convincing arguments against P1. But it would then be these arguments that would in fact do all the work, since blocking P1 would already by itself block the argument to C1.

One might still argue that the claim that the (primary) intension of the word 'crimson' includes certain superficial looks (under normal circumstances) is true regarding most competent users of the word, by endorsing a subject-sensitive notion of (primary) intension. Jackson & Chalmers are indeed sympathetic towards such a notion:

"[P]erhaps a city-dweller might use 'water' nondeferentially for the liquid that comes out of faucets (knowing nothing of oceans), and a beach-dweller might use 'water' nondeferentially for the liquid in the oceans (knowing nothing of faucets). This sort of case can be treated as we suggested that cases involving names be treated: the subjects have different conditional abilities, and different associated conditionals will be knowable a priori for each of them." (Jackson & Chalmers 2001, p. 328)

However, resorting to a subject-sensitive notion of (primary) intension does not help in order to block the argument to C1. For C1 is merely a claim about the *modified* version of the case of Norman, and employing a subject-sensitive notion of (primary) intension does nothing to threaten this conclusion. In fact, it might even make it more plausible,

for given a subject-sensitive notion of (primary) intension, it seems quite plausible to assume that the (primary) intension of the word ‘crimson’, as used by Norman in the modified version of the scenario, does not include any particular superficial looks.

I conclude that both ways of explaining the required classification of certain ways of understanding as primary (that is, explaining it on the basis of an undifferentiated notion of intension as well as explaining it on the basis of the notion of primary intension) are unacceptable within a strategy for blocking the argument to C1 that does not provide any independent reasons for rejecting P1. Since the two-dimensionalist way of explaining the required classification appears to be the most modest way available of grounding this classification in the contents of the respective expressions, and since it is not at all clear how it could plausibly be denied that primary and secondary ways of understanding of words are reflected in the words’ contents, it can also be concluded that the strategy of blocking the argument to C1 on the basis of a version of NE is unsuccessful without providing independent support for P1.

2.5. Summary and outlook

In the previous sections, I have considered a possible attempt at blocking the argument to C1 on the basis of the notion of degrees of understanding. As described in section 2.2, in order for this strategy not to overgeneralize, it has to resort to a classification of certain ways of understanding a word as primary and of certain other ways of understanding a word as secondary. In section 2.3, I have considered two possible ways of grounding this classification in the contents of the respective words. It has turned out that the more modest of these two attempts, which employs a distinction between different kinds of intensions, is the more promising one. In section 2.4, however, I have argued that both attempts fail, since they turn out to be circular within a strategy for blocking the argument to C1 that does not involve any independent argument against P1. Given that it is not at all clear how it could be convincingly denied that the required classification of certain ways of understanding a word is reflected in a word’s content, and given that the more modest of the two attempts appears to be the most modest way available of grounding the required classification in the contents of the respective expressions, I have concluded that the strategy of blocking the argument to C1 on the basis of a version of NE ultimately fails.

At the end of the previous section, I have pointed out that endorsing a subject-sensitive notion of (primary) intension does not help in order to block the argument to C1. However, when it comes to evaluating the plausibility of P3, the prospects for defending such a notion are quite relevant. This issue will therefore be dealt with in a later chapter (namely, in section 4.3).

Since, as I have argued in section 2.4, the considered ways of grounding the required classification of ways of understanding turn out to be circular without providing independent support for P1, I will now, in the following chapter, consider an argument that might provide such support.

3. Examining the argument from analogy, part II: P1 and the primary objects of truth

In this chapter, I will consider an argument against P1 that mainly relies on a certain view about the fineness of grain of the primary objects of (a priori) truth. In section 3.1, I will develop this argument and explain in how far it would, if successful, undermine P1. In section 3.2, I will present some *prima facie* considerations in favour of a view about the fineness of grain of the primary objects of (a priori) truth whose correctness would constitute the unsoundness of the argument against P1. In sections 3.3 and 3.4, I will argue that certain worries concerning this view do not outweigh the *prima facie* considerations in favour of it presented in section 3.2. Finally, in section 3.5, I will argue that my case for the unsoundness of the argument against P1 also constitutes a positive case in favour of P1.

3.1. The argument against P1

So far, my discussion of the epistemic status of Norman's knowledge of R1 has been quite unspecific concerning the question of what kind of entity R1 should be regarded as. That is, I have remained unspecific concerning the question of whether R1, if it is to be regarded as an entity of a kind which is the primary object of a priori truth (and thereby of a kind which is the primary object of truth), should thereby be regarded as a linguistic entity (that is, as an interpreted or uninterpreted sentence), or as an entity that would usually be called a proposition. Instead, I have merely introduced R1 and R2 as truths, thereby remaining as neutral as possible regarding the questions just mentioned. As it will occur in the present chapter, these questions are highly relevant regarding the prospects of the argument to C1, and in particular regarding the role played by P1 in this argument – that is, regarding the role played by the claim that in virtue of possessing the skills described in the example of Mary and Norman, Norman understands the word 'crimson'.

In order to see the relevance of the above-mentioned questions for the plausibility of P1, it is helpful to provide some specification of the claim P1 is supposed to consist in. Assuming, for example, that R1, if it is to be regarded as an entity of a kind which is

the primary object of truth, should thereby be regarded as a relatively fine-grained non-linguistic entity, with concepts (individuated through cognitive criteria) as its constituents, then the truth of P1 will only be relevant regarding the epistemological issues at hand if it involves the claim that in virtue of possessing the skills described in the example of Mary and Norman, Norman grasps the relevant colour concept that figures in P1. If R1 should be regarded as a relatively fine-grained entity of the kind just mentioned, and if P1 did *not* directly or indirectly involve the claim just mentioned, then P1 would be rather irrelevant regarding the epistemological status of R1. Defenders of the argument from analogy should therefore be prepared to accept an interpretation of P1 along the following lines:

(P1_{INT}) In virtue of possessing the reliable skill of correctly applying the word ‘crimson’ on the basis of observing the surfaces of crimson objects through a microscope, Norman understands or grasps the constituent of R1 contributed by the word ‘crimson’.

This interpretation of P1 is not supposed to be inconsistent with the possible claim that R1, if it is to be regarded as an entity of a kind which is the primary object of truth, should thereby be regarded as a sentence. In this case, the connection between understanding the word ‘crimson’ and understanding or grasping the corresponding constituent of R1 would simply be a very direct one, since the word ‘crimson’ would then be identical to the corresponding constituent of R1.

Given this interpretation of P1, it now becomes more visible why the above-mentioned questions are relevant regarding the truth of P1. A case particularly threatening for defenders of P1 would be one in which it turns out that the primary objects of truth are cognitively individuated, and thereby relatively fine-grained, non-linguistic entities, such that R1 should be regarded as more fine-grained than the sentence that I have hitherto used in order to express it – that is, as more fine-grained than the sentence ‘All crimson things are red’. In this case, a relatively extensive amount of interpretation would lie between R1 and the sentence just mentioned. That is, if the constituents of R1 are cognitively individuated entities, then it will, for example, be a matter of interpretation which entity associated with the word ‘crimson’ figures in R1.

The argument against P1 that I will consider in this chapter¹⁹ rests on the assumption that the kind of entity exemplified by R1, if R1 is to be regarded as the kind of entity which is the primary object of truth, is relatively fine-grained, such that it is indeed a matter of interpretation which entity associated with the word ‘crimson’ figures in R1.²⁰ One possible way of specifying this assumption involves the idea of different concepts of crimson, only one of which figures in R1. Defenders of the apriority of R1 who wish to refute P1 might then go on to claim that R1 – as they have always conceived of it, and as it should therefore be interpreted if the argument from analogy is to be relevant at all regarding the debate in question – contains as one of its constituents a concept of crimson (and *only* this concept of crimson) that Norman clearly does not grasp in the scenario involving Mary and him. That is, defenders of the apriority of R1 might claim that what figures in R1 is a *phenomenal* concept of crimson – one that Norman clearly does not acquire by being taught about microphysical surface properties. It could therefore be concluded that, if the argument from analogy is to be relevant at all regarding the debate in question – that is, if it is to be directed against a claim regarding colours that main defenders of the a priori have actually made – then P1 is false.

As stipulated regarding the scenario of Mary and Norman, Norman has seen crimson objects (under good lighting conditions), and since his visual apparatus functions normally, he is thereby acquainted with the phenomenal character of crimson. However, it might well be claimed that, since Norman does not associate this phenomenal character with the term ‘crimson’, or with any other term, the boundaries of this phenomenal character do not figure in whatever might plausibly be called his colour categorization scheme. And from this, it might well be concluded that Norman does not grasp the phenomenal concept of crimson. But even if Norman’s acquaintance with the phenomenal character of crimson already constitutes a grasp of the respective phenomenal concept, it is still safe to say that, contrary to what is stated in P1, Norman does not possess the relevant grasp or understanding in virtue of his skills related to microphysical surface properties.

¹⁹ It is an argument that was suggested to me by comments on chapter 1 made by Laurence Bonjour via personal communication.

²⁰ An idea along such lines can be found in Peter Achinstein’s 1968 book *Concepts of Science*. Achinstein claims that the sentence ‘Copper is the element of atomic number 29’ can “be used to express two different *statements*” (p. 41), one being tautological, and the other being decidable only by empirical means – due to contributions of the term ‘copper’ differing regarding degrees of superficiality. Nothing in Achinstein’s remarks suggests that he would propose a different kind of treatment regarding sentences containing the term ‘crimson’.

I call the foregoing argument the argument from fineness of grain. In order to facilitate evaluation of the argument, here is a list of its premises (presenting the argument in a slightly more charitable way than in the above description):

- (P_{FG1}) Whatever concept Norman associates with the word ‘crimson’²¹, it significantly differs from the concept which people who have acquired the term ‘crimson’ in the usual way associate with the word ‘crimson’.
- (P_{FG2}) The primary objects of a priori truth (and thereby the primary objects of truth) are fine-grained enough to be sensitive to the differences between concepts.
- (P_{FG3}) If both P_{FG1} and P_{FG2} are true, then an interpretation of the sentence ‘All crimson things are red’ relevant to what is at issue in debates about apriority should be one according to which the entity contributed to R1 by the term ‘crimson’ is directly connected to the phenomenal character of this colour.

My goal in the next few sections will be to refute the argument from fineness of grain. My strategy will be the following. Since it is hard to deny that it follows from the above-listed premises that P1 – if it is to be relevant to what is at issue in debates about apriority – is false, I will try to refute the argument from fineness of grain by arguing that it is unsound. P_{FG1} might be challenged on the basis of a view about concepts according to which concepts are not individuated cognitively, but in a more coarse-grained way. However, even a successful argument along these lines would still leave available to defenders of the argument from fineness of grain the possibility to resort to a modified version of P_{FG2} according to which the primary objects of truth are not only fine-grained enough to be sensitive to the differences between concepts, but also fine-grained enough to be sensitive to the differences between certain more fine-grained, cognitively individuated entities. Accordingly, my strategy for refuting the argument from fineness of grain will consist in attempting to refute the claim that the primary objects of truth are fine-grained enough to be sensitive to the differences between cognitively individuated entities, regardless of whether concepts are individuated in this

²¹ The aim behind this formulation is to remain neutral regarding whether the concept that Norman associates with the term ‘crimson’ deserves to be called a concept *of crimson*. Presenting the argument against P1 as presupposing that Norman possesses a concept of crimson would be less charitable, since it would, for example, make P_{FG2} more vulnerable.

way. That is, I will argue that, if P_{FG1} is false, then – even though P_{FG2} might be true in this case – the modified version of P_{FG2} just mentioned will be false; and if P_{FG1} is true, then not only the modified, but also the original version of P_{FG2} will be false. If my strategy is successful, then P_{FG3} will be vacuously true. But even though P_{FG3} 's being *non-vacuously* true would by itself establish the truth of the conclusion of the argument from fineness of grain, P_{FG3} 's being vacuously true would be of no help for defenders of the argument.

3.2. The coarse-grained view and some initial considerations in favour of it

As already stated in the preceding section, I will try to refute the argument from fineness of grain by arguing that the primary objects of truth are not fine-grained enough to be sensitive to certain differences between cognitively individuated entities. In particular, I will try to argue that whatever cognitive differences occur between Norman's way of thinking about what he calls 'crimson' and the way in which people who learned the word 'crimson' in the usual way think about this colour, these differences are none that the primary objects of truth are sensitive to.

As already mentioned in the introduction to the dissertation, my basic assumption that apriority can felicitously be ascribed to the objects of truth does not carry with it any assumption regarding whether the primary objects of truth are propositions, sentences, etc. In my argument that the primary objects of truth are not sensitive to the relevant differences between Norman and the others, I will continue to stay neutral regarding this issue. It should be noted that those regarding sentences as the primary objects of truth should immediately reject the argument from fineness of grain, since the view that sentences are the primary objects of truth is clearly incompatible with the conjunction of P_{FG1} and P_{FG2} . However, even though Jackson, for instance, might be interpreted as endorsing the view just mentioned by stating that "by an a priori (true) sentence I will mean one such that understanding it is sufficient for being able to see that it is true" (2000, p. 324; see also his 1998, p. 52, and his 2004, p. 258), the view should be regarded as a non-standard one.

The only initial assumption I make regarding the primary objects of truth is the assumption that the primary objects of truth are the primary objects of knowledge. Given this assumption, which I take to be uncontroversial, one can now provide a

clarification of the claim that the cognitive differences between Norman's way of thinking about what he calls 'crimson' and the way in which people who learned the word 'crimson' in the usual way think about this colour are none that the primary objects of truth are sensitive to. That is, the claim can now be reformulated as one about sameness of what is known in certain situations. More exactly, it can be reformulated as follows: If Norman knows a sentence containing the word 'crimson' to be true (or to express a truth), and people who learned the word 'crimson' in the usual way also know the respective sentence to be true (or to express a truth), then (assuming that the example does not involve homonymy, as in 'This bank is crimson') what is known by Norman and by the other people in this case is the same, regardless of the cognitive differences between their ways of thinking about what they call 'crimson'.

In what follows, I will, for practical purposes, use the expression 'fine-grained view' to cover all views about the objects of knowledge according to which what is known in the instances of knowledge just described is *not* the same, due to the cognitive differences in question. Correspondingly, I will use the expression 'coarse-grained view' to cover all views about the objects of knowledge according to which the same is known in the described instances of knowledge, regardless of the cognitive differences in question. Of course, proponents of the coarse-grained view should not deny that the cognitive differences in question result in differences between what is known (or at least between what is believed) by Norman and the others. The coarse-grained view merely consists in the claim that in the particular instance of knowledge corresponding to the above-mentioned sentence, which is known to be true (or to express a truth) both by Norman and the others, the objects of knowledge are the same. In other words: In the particular instance of knowledge corresponding to the above-mentioned sentence, the word 'crimson' does not contribute different entities to what is known by Norman and to what is known by the others.

My strategy for refuting the argument from fineness of grain can now simply be redescribed as the attempt at making a case for the coarse-grained view. In the remainder of the present section, I will present some *prima facie* considerations favouring the coarse-grained view. These considerations do certainly not constitute a definite case for the view. Rather, they are merely supposed to shift the burden of proof to proponents of the argument from fineness of grain.

The considerations I have in mind simply consist in an appeal to (presumably widespread) intuitions regarding sameness of what is known in some example cases. Since

the coarse-grained view is primarily concerned with the difference between Norman and people who learned the word ‘crimson’ in the usual way, I will, in order for the considerations in question to be as relevant as possible, start with two examples involving this difference. Consider, first, the following sentence:

(S1) ‘Crimson is a colour.’

Unsurprisingly, both Norman and the others sincerely assent to this sentence. According to the fine-grained view, the corresponding beliefs of Norman and the others have different objects, since, due to the difference between their respective ways of thinking about what they call ‘crimson’, the word ‘crimson’ contributes – according to the fine-grained view – different entities to what is believed by Norman and the others. However, from an intuitive point of view, it is quite unproblematic to interpret both Norman’s and the others’ assent to S1 via a single attitude report, that is, to ascribe to both parties the same belief, namely, the belief that crimson is a colour. Given that the example is one in which neither Norman nor the others are faced with any misleading evidence or with any unusual circumstances reducing the stability of their sincere and clear-headed assent to S1, it is safe to say that their belief amounts to knowledge. Denying that Norman’s belief amounts to knowledge would commit one to denying the possibility of knowledge through testimony, since this is how Norman forms the belief in question; and the example of Mary and Norman might be construed so as to contain an arbitrarily high number of maximally reliable people telling Norman that the word ‘crimson’ is a colour term. Thus, from an intuitive point of view, the example involving S1 provides support for the coarse-grained view. Consider another sentence:

(S2) ‘This book is crimson.’

Suppose that this time, Norman’s assent is not based on testimony, but on an observation of the book in question through a microscope. That is, in contrast to the example involving S1, Norman’s coming to assent to the sentence in question depends directly on his particular way of thinking about what he calls ‘crimson’. But nevertheless, from an intuitive point of view, it is once again unproblematic to ascribe the same belief to Norman and the others, namely, the belief that this book is crimson.

And given that the present example is construed analogously to the preceding one, it is again safe to say that both Norman's and the others' beliefs amount to knowledge.

It is plausible to suppose that the above observations regarding the intuitive appeal of the coarse-grained view generalize to all other cases that the view is directly concerned with, that is, to all cases in which Norman knows a sentence containing the word 'crimson' to be true (or to express a truth), and in which people who learned the word 'crimson' in the usual way also know the respective sentence to be true (or to express a truth) – given, of course, that no kind of ambiguity in the rest of the sentence plays any role. The next question to ask is how much weight should be given to the fact that the coarse-grained view is intuitively appealing when applied to the above kind of example. In particular, it might be pointed out that the kind of intuition at issue here – namely, intuitions regarding sameness of what is believed or known – will only qualify as reliable if they cannot be produced too easily. But such worries turn out to be unnecessary, since the kind of intuition at issue already vanishes when it comes to examining sentence pairs generated by a substitution of terms that are coextensive with respect to every possible world. Just consider the following classic example:

(S3) 'Hesperus is Hesperus.'

(S4) 'Hesperus is Phosphorus.'

From an intuitive point of view, the corresponding knowledge reports (yielded by disquotation) ascribe different pieces of knowledge to the respective subjects. Of course, so-called intuitions regarding these kinds of issues may change in the course of one's theoretical engagement with them. However, regarding the present example, the intuitive appeal of Frege's initial claim that " $a=a$ and $a=b$ are obviously statements of differing cognitive value" (Frege 1948, p. 209) seems to be wide-spread and relatively stable, such that proponents of Millianism – the view according to which coextensive proper names contribute the same entity to the objects of knowledge – found themselves faced with the task to explain away these intuitions. Nathan Salmon, for example, in his classic defense of the Millian account, explicitly states that this account, applied to the present example, commits one to the claim that "anyone who knows that Hesperus is Hesperus knows that Hesperus is Phosphorus, no matter how strongly he or she may deny the latter." (Salmon 1986, p. 83) Admitting that claims like this "clash sharply

with ordinary usage” (p. 84), Salmon offers a pragmatic account of attitude reports in order to explain away the relevant intuitions (see esp. pp. 114ff.).

Theories of propositions according to which the latter are sets of possible worlds and also the objects of the attitudes (see Stalnaker 1976 and 1984, chs. 3 and 4, for a classic statement and defense) will of course be faced with even more stable recalcitrant intuitions. (See Soames 1987, pp. 47-50, for some standard criticism.) It will not be necessary for present purposes to go into more detail here. The above observations already suffice to show that the intuitions regarding examples like S1 and S2 are not inflationary in character. That is, these intuitions do not overgeneralize so as to support just *any* degree of coarseness of grain. The fact that there is some threshold of intuitiveness somewhere between Millianism and a maximally fine-grained version of what I have called the coarse-grained view does of course provide Millianism, as well as even more coarse-grained views, with some prima facie implausibility, but it also provides a version of the coarse-grained view which is fine-grained enough to rule out Millianism with some prima facie plausibility. Moreover, I take it that the above-described intuitions, taken together, provide the fine-grained-version of the coarse-grained view just described (that is, the view according to which what I have called the coarse-grained view is true, but according to which Millianism is still false) with some prima facie plausibility favouring it over all versions of what I have called the fine-grained view. Thus, I also take it that the burden of proof lies with those who defend the fine-grained view, that is, the view that the term ‘crimson’, as used by Norman, contributes a different entity to the objects of knowledge than the same term, as used by people who have learned it in the usual way.

Before examining, in the next two sections, whether the prima facie plausibility of the coarse-grained view might be defeated by certain objections, it is important to emphasize again that the above-described prima facie considerations regarding intuitions about sameness of what is known do not only render prima facie implausible those theories according to which the term ‘crimson’, as used by Norman, contributes a different entity to the *proposition* expressed than the same term, as used by people who have learned it in the usual way. Rather, the above-described considerations render prima facie implausible all theories according to which different entities are contributed to the *objects of knowledge*. Among those theories is a kind of view according to which the term ‘crimson’ contributes the same entity to the proposition expressed, but according to which the difference between the ways Norman and the others think about

what they call ‘crimson’ nonetheless affects the truth-conditions of the corresponding attitude reports. This is because – according to the kind of view I have in mind – attitudes like belief and knowledge are not, as traditionally assumed, two-place relations between a subject and a proposition, but instead three-place relations involving, as a third relatum, a way of grasping (or understanding, thinking of, etc.) the respective proposition. This third relatum may be partly determined by the choice of words in the formulation of the attitude reports, as in ‘S knows that Hesperus is Venus’ vs. ‘S knows that Phosphorus is Venus.’ Concerning report pairs like ‘Norman knows that this book is crimson’ vs. ‘The others know that this book is crimson’, defenders of the view in question will hold that the third relata are implicitly determined by the respective contexts of utterance. (Versions of such an account are, for example, developed in Schiffer 1977, Fitch 1984, Crimmins & Perry 1989, and Crimmins 1992.) What is important to note, for present purposes, regarding this kind of contextualist view is that it should plausibly be interpreted as one according to which the primary objects of attitudes like belief and knowledge are not just propositions, but propositions grasped (understood, thought of, etc.) in a certain way. Therefore, the above-described intuitions regarding S1 and S2 are not just directed against views according to which propositions (conceived of as the primary objects of the attitudes) are relatively fine-grained, but also against the contextualist view in question; for also according to the latter, Norman’s assent to S1 and S2 amounts to knowledge of different entities than the others’ assent to S1 and S2, due to a difference between Norman’s and the others’ way of grasping the propositions in question.

Moreover, even if the contextualist view in question should (implausibly) be interpreted as one according to which the primary objects of attitudes like belief and knowledge are just propositions, and would thereby (given the corresponding coarse-grained conception of the primary objects of knowledge) not be in conflict with the above-described intuitions regarding S1 and S2, it would still involve counter-intuitive consequences regarding the case of Norman and the others. Consider possible contexts of utterance in which, according to the view, the belief report ‘John (who learned the term ‘crimson’ in the usual way) knows that this book is crimson’ expresses a truth. These would be contexts in which a certain way of grasping (understanding, thinking of, etc.) the proposition that this book is crimson is salient in a certain way – namely, a way of grasping (understanding, thinking of, etc.) related to the superficial looks of crimson things. Accordingly, the view predicts that, in such contexts, uttering the belief report

‘Norman knows that this book is crimson’ expresses a falsehood. But it seems, at least intuitively, wrong to accept that contextual features regarding salience of a certain way of grasping (understanding, thinking of, etc.) could have this kind of impact. That is, no matter how salient a way of grasping (understanding, thinking of, etc.) related to the superficial looks of crimson things might be, the belief report ‘Norman knows that this book is crimson’ would still seem to express a truth. (A similar line of argument can be found in Schiffer 1994, p. 287.) I therefore take it to be at least *prima facie* plausible to assume that, if a version of what I have called the coarse-grained view is correct, then it will not be the contextualist view just considered.

My main aim in the present section has been to show that certain *prima facie* considerations prefer a certain (in the relevant sense non-contextualist) fine-grained version of what I have called the coarse-grained view over what I have called the fine-grained view. These considerations thereby shift the burden of proof to proponents of the argument from fineness of grain. In the following two sections, I will examine whether certain worries regarding the fine-grained version of the coarse-grained view can outweigh the *prima facie* plausibility in question.

3.3. Translation

One might wonder what a theory of the constituents of propositions that is in accordance with a combination of the fine-grained version of the coarse-grained view and the popular view that propositions are the objects of the attitudes might look like. Such a theory of the constituents of propositions is suggested (though not fully endorsed) in Bealer 1998. According to the proposal in question, what names contribute to the proposition expressed are linguistic entities of a certain kind. Bealer considers different possibilities of conceiving of the kind of linguistic entity in question (corresponding to the labels ‘living names’, ‘naming practices’, and ‘historical naming trees’), the details of which are irrelevant for present purposes. What is essential is that the kind of entity in question is, in some way or other, determined by the combination of a word and its referent (cf. Bealer 1998, pp. 16f.). The referent’s being involved in the kind of linguistic entity in question ensures that homophonous names do not always contribute the same entity to the proposition expressed – obviously a necessary feature of the theory at hand in order to be promising. The *word*’s being involved in the kind of

linguistic entity in question ensures that – contra Millianism – coextensive but orthographically distinct names contribute different entities to the respective propositions. Thus, regarding the sentence pair ‘Cicero is a person’ and ‘Tully is a person’, Bealer, conceiving of the propositions expressed as what he calls ‘descriptive predications pred_d ’, holds that “[g]iven that ‘Cicero’ and ‘Tully’ [i.e., the linguistic entities contributed by the names ‘Cicero’ and ‘Tully’] are distinct, $\text{pred}_d(\text{being a person, ‘Cicero’})$ and $\text{pred}_d(\text{being a person, ‘Tully’})$ would be distinct” (p. 18).

At that point, Bealer presents the account as one according to which only names contribute the kind of linguistic entity in question to the proposition expressed, while predicates, like ‘is a person’, merely contribute their referent, that is, the corresponding property. It is, however, not at all clear why name pairs like ‘Cicero’ and ‘Tully’ should be treated differently from pairs of predicates that are coextensive with respect to every possible world, like ‘being water’ and ‘being H_2O ’. Given that, as is widely acknowledged, ‘being water’ and ‘being H_2O ’ denote the same property, claiming that these terms only contribute their referents would commit one to accepting counter-intuitive consequences analogous to those Bealer wishes to avoid by treating names in the way described above. Accordingly, towards the end of his paper, in a section called “Applications”, he proposes to treat the difference between synonymous predicates, like ‘chew’ and ‘masticate’, analogously to the difference between coextensive, but orthographically distinct, names (cf. p. 28). Obviously, this proposal automatically extends to one according to which *all* predicates (and not just certain pairs of synonymous ones), including ‘being crimson’, contribute the kind of linguistic entity described above to the proposition expressed.

As already indicated above, Bealer does not fully endorse the account just described. More exactly, he holds that, instead of linguistic entities involving both a word and its referent, what is contributed to the proposition expressed may also be construed as “e.g. clusters of recognitional routines causally involving Cicero in an essential way; percepts which are essentially individuated by their objects; expressions in the “language of thought”; etc.” (p. 20), and he notes that “[i]n the present context, there is no need to declare any one of these proposals to be most promising” (ibid.). Accordingly, my aim in presenting the above account was not to provide substantial support for the fine-grained version of the coarse-grained view, but merely to show what kind of theory of the constituents of propositions a combination of the fine-grained version of the coarse-grained view and the popular view that propositions are the

objects of the attitudes might result in. I take it that the above description of Bealer's proposal has shown that such a theory of the constituents of propositions can be construed quite easily. And even though this fact does not provide any substantial support for the fine-grained version of the coarse-grained view, it does certainly not diminish its *prima facie* plausibility.

However, demonstrating what kind of theory of the objects of the attitudes is suggested by the fine-grained version of the coarse-grained view reveals that the view is apparently in tension with common practices of translating attitude reports from one language into another and of reporting attitudes of people in a language that these people do not understand. According to the above-described account of the objects of the attitudes, the belief report 'Some ancient Greeks believed that the Earth is round' expresses a falsehood, simply because none of the ancient Greeks understood modern English, and therefore, so one might plausibly assume, none of them stood in a belief relation to the entity expressed by the *that*-clause of the belief report in question, if the constituents of this entity involve the respective words. Commitments of this kind may be taken as an unwelcome consequence of endorsing the fine-grained version of the coarse-grained view.

Moreover, proponents of the fine-grained view can quite easily avoid the unwelcome commitment in question. In order to do so, they would, of course, have to endorse a version of the view according to which the constituents of the objects of the attitudes do *not* involve words or similar linguistic entities. Instead, they might opt for one of the alternatives already suggested by Bealer, e.g., clusters of recognitional routines causally involving the object in question in an essential way, or percepts essentially individuated by their objects. The corresponding theories would qualify as versions of what I have called the fine-grained view, since it seems plausible to hold that Norman and the others sufficiently differ regarding their recognitional routines or percepts associated with the word 'crimson' in order for the objects of their attitudes corresponding to sentences containing the word 'crimson' to differ, given that the kind of theory in question is correct. And claiming there to be this difference between Norman and the others still seems to leave room for claiming that many ancient Greeks and many modern-day speakers of English are sufficiently similar regarding their recognitional routines or percepts associated with the referents of any of the words contained in 'The Earth is round' in order for some of the ancient Greeks to stand in a belief relation to the same entity as most modern-day speakers of English assenting to

the sentence in question. It seems, therefore, that theories of the objects of the attitudes (more or less) suggested by the fine-grained view fare clearly better than those suggested by the fine-grained version of the coarse-grained view when it comes to issues regarding translation. And this fact might well be taken to outweigh the *prima facie* considerations described in the previous section.

Such a conclusion, however, would be too hasty. As already mentioned, proponents of the fine-grained view can only avoid the unwelcome consequences in question by opting for a version of their view according to which the constituents of the objects of the attitudes do not involve words or similar linguistic entities. But endorsing such a version of the fine-grained view involves its own unwelcome consequences. Consider the words ‘chew’ and ‘masticate’. Since these words stand in a relation to each other that is commonly called a relation of synonymy, it seems plausible to hold that in the case of most competent speakers regarding the two words, the recognitional routines or percepts associated with the word ‘chew’ do not significantly differ from those associated with the word ‘masticate’. At least, proponents of the version of the fine-grained view in question can hardly coherently hold that, on the one hand, the recognitional routines or percepts associated with the word ‘Earth’ by some modern-day speakers of English do not significantly differ from those figuring in beliefs of some ancient Greeks, and that, on the other hand, the recognitional routines or percepts associated with the words ‘chew’ and ‘masticate’ by a single modern-day speaker of English do not in general coincide.

Thus, proponents of the version of the fine-grained view in question, if they want to avoid the unwelcome consequences regarding translation, would have to hold that substituting ‘chew’ for ‘masticate’, or vice versa, in belief reports uttered by competent speakers regarding these words does in general not result in the attribution of a different belief. As a consequence, proponents of the version of the fine-grained view in question are committed to accepting that these speakers attribute one and the same belief when uttering sentences of the form ‘x believes that whoever chews chews’ and ‘x believes that whoever chews masticates’. The version of the fine-grained view in question thereby turns out to be more coarse-grained regarding certain cases than the fine-grained version of the coarse-grained view. In order to see that the kind of consequence just mentioned should be regarded as unwelcome, just consider a person S who knows the meaning of the word ‘chew’, but has no idea what the word ‘masticate’ means. Given the consequences already mentioned of endorsing the version of the fine-grained view

in question, proponents of the latter are committed to claiming that most competent speakers regarding the two words would either express two truths or two falsehoods when uttering ‘S believes that whoever chews chews’ and ‘S believes that whoever chews masticates’. I take this commitment to be clearly unwelcome, since it is quite intuitive to hold that only the first, but not the second of these two utterances expresses a truth.

Thus, while theories of the objects of the attitudes suggested by the fine-grained version of the coarse-grained view involve counter-intuitive consequences regarding common practices of translation, versions of the fine-grained view that avoid these consequences are faced with counter-intuitive commitments of their own. Instead of examining how proponents of the fine-grained version of the coarse-grained view might try to avoid the counter-intuitive consequences that proponents of Bealer’s proposal are faced with, it shall, for present purposes, suffice to point out that proponents of the fine-grained view, if they wish to avoid these consequences, are faced with commitments that do not seem any less unwelcome than those that proponents of Bealer’s proposal are faced with. One can thereby already conclude that considerations regarding translation do not outweigh the prima facie considerations described in the previous section. At least, the burden of proof now lies with those who claim otherwise.

In the following section, I will consider a second kind of worry that might be regarded as outweighing the prima facie considerations described in section 3.2.

3.4. Kripke’s Paderewski case

When examining the plausibility of what I have called the coarse-grained view – the view that the primary objects of knowledge are not sensitive to the kind of difference between Norman’s way of thinking about what he calls ‘crimson’ and the way in which people who learned the word ‘crimson’ in the usual way think about this colour – it is necessary to consider a certain kind of scenario, which has been introduced into the philosophical literature by Saul Kripke in his 1979 paper “A Puzzle about Belief”. (Kripke’s own motivation for introducing the scenario shall be left aside for the moment.) The central feature of the kind of scenario in question is that it involves a subject who is acquainted with different features of a certain object, not being aware that these different features are in fact features of one and the same object. Moreover,

the subject uses the same name for what she takes to be two different objects. That is, she uses the name in question assuming the case to be one of homophony, while in fact there is only one object referred to.

In the particular case described by Kripke, the object in question is a person named Paderewski (actually a real world example), who was both a famous pianist and a famous Polish statesman. Peter – the subject in the example – comes to associate the name ‘Paderewski’ with both of these features, but he is unaware that the name refers to only one person. Corresponding to this unawareness, Peter comes to assent to (at least superficially) contradictory sentences:

“Peter [...] may learn the name ‘Paderewski’ with an identification of the person named as a famous pianist. Naturally, having learned this, Peter will assent to “Paderewski had musical talent,” and we can infer – using ‘Paderewski,’ as we usually do, to name the Polish musician and statesman:

(8) Peter believes that Paderewski had musical talent.

[...] Later, in a different circle, Peter learns of someone called ‘Paderewski’ who was a Polish nationalist leader and Prime Minister. Peter is skeptical of the musical abilities of politicians. He concludes that probably two people, approximate contemporaries no doubt, were both named ‘Paderewski.’ Using ‘Paderewski’ as a name for the *statesman*, Peter assents to, “Paderewski had no musical talent.” (Kripke 1979, p. 265)

As Kripke points out here, inferring from Peter’s assent to ‘Paderewski had musical talent’ that he believes (at least before he also becomes acquainted with the name ‘Paderewski’ in the other way mentioned) that Paderewski had musical talent corresponds to people’s usual practice of belief attribution. The central question concerning the scenario is whether an analogous belief report should be made regarding Peter’s later assent to ‘Paderewski had no musical talent’. Doing so, while assuming that Peter has not given up the belief that Paderewski had musical talent, would result in attributing to him two (at least superficially) contradictory beliefs held at the same time.

Kripke’s evaluation of this scenario can be extracted from his remarks concerning his well-known example of Pierre, who is unaware that the names ‘Londres’ and ‘London’ are names of the same city (see pp. 254ff.). This case differs from the one involving Peter basically only insofar as the derivation of attributions of (at least superficially) contradictory beliefs in the case of Pierre involves a principle according to which sincere assent to a sentence in a certain language (in this case the French sentence ‘Londres est jolie’) can correctly be interpreted via a belief report formulated in another language (in this case the English belief report ‘Pierre believes that London is pretty’). Kripke’s description of the Paderewski case, which directly follows his evaluation of

the case involving Pierre, has the mere function of demonstrating that rejecting such a translation principle is of no help in solving the kind of problem in question (cf. p. 266). Since the problem involving Peter and Paderewski is otherwise supposed to be “parallel to the problem with Pierre and London” (ibid.), Kripke’s evaluation of the former is supposed to be inferred from his evaluation of the latter.

As already mentioned, attributing to Peter the belief that Paderewski had no musical talent, while holding that he has not given up any of his beliefs about the person he is acquainted with as Paderewski-the-pianist, would result in attributing (at least superficially) contradictory beliefs to Peter. Even though Kripke does not argue against the general possibility of someone’s holding contradictory beliefs at the same time, he claims that, in the example at hand, making such attributions would be subject to “insuperable difficulties” (p. 257). Regarding the case of Pierre and London, which Kripke, as mentioned above, evaluates analogously to that of Peter and Paderewski, one might simply stipulate that Pierre is a leading logician, who – as Kripke assumes – “would *never* let contradictory beliefs pass” (ibid.). And even without this kind of stipulation, Kripke regards the belief attributions in question as untenable:

“[S]urely anyone, leading logician or no, is in principle in a position to notice and correct contradictory beliefs if he has them. Precisely for this reason, we regard individuals who contradict themselves as subject to greater censure than those who merely have false beliefs. But it is clear that Pierre, as long as he is unaware that the cities he calls 'London' and '*Londres*' are one and the same, is in no position to see, by logic alone, that at least one of his beliefs must be false. He lacks information, not logical acumen. He cannot be convicted of inconsistency: to do so is incorrect.” (Kripke 1979, p. 257)

Obviously, the only way to avoid this (in Kripke’s view) undesirable result is to refrain from attributing to Peter at least one of the two beliefs in question. According to Kripke, however, this option is unattractive as well. Given that it seems very hard to deny that Peter at least once had the belief that Paderewski had musical talent, claiming that he does not have this belief after learning about Paderewski-the-statesman would commit one to claiming that Peter *lost* the belief that Paderewski had musical talent, due to which, in turn, according to Kripke, “we would be forced to say that Pierre [and thereby also Peter] has *changed his mind, has given up his previous belief*” (p. 256). And this, Kripke holds, would be incorrect, as soon as it is stipulated that the subject (be it Pierre or Peter) “says he has not changed his mind about anything, has *not* given up any belief” (ibid.).

Regarding the option of refraining from attributing to Peter the later belief that Paderewski had no musical talent, Kripke is similarly sceptical. In analogy to a possible extension, considered by Kripke (cf. p. 257), of the example of Pierre and London, one might consider an extension of the example of Peter in which an electric shock wiped out all his memories of what he learned about Paderewski-the-pianist, leaving Peter with exactly the kind of information about Paderewski that many other people who only know him as a statesman also possess – people to whom (given that they, like Peter, believe that politicians have no musical talent) one should clearly attribute the belief that Paderewski had no musical talent. Kripke now holds that “surely no shock that *destroys* part of Pierre's [or Peter's] memories and knowledge can *give* him a new belief” (ibid.). So since there seems to be no good reason not to attribute to Peter the belief in question after the electric shock, and since, according to Kripke, the electric shock could not have created the belief, it seems unavoidable to accept that Peter believes that Paderewski had no musical talent, regardless of whether he still has his memories about Paderewski-the-pianist.

Kripke's diagnosis regarding the kind of scenario at hand is that it presents a genuine puzzle, one that “indicates that the usual principles we use to ascribe beliefs are apt, in certain cases, to lead to contradiction, or at least, patent falsehoods” (p. 267). His motivation for presenting the puzzle is the following. According to Millianism about names, coextensive proper names contribute the same entity to the objects of belief and knowledge. The Millian is thereby committed to the claim that coextensive proper names can always be substituted *salva veritate* within attitude reports – which, as already seen earlier, yields counter-intuitive results. Kripke's aim in presenting the scenarios of Pierre and London and of Peter and Paderewski is to suggest that rejecting Millianism on the basis of these counter-intuitive results would be too hasty, since, as the cases in question are supposed to demonstrate,

“the absurdities that disquotation plus substitutivity would generate are exactly paralleled by absurdities generated by disquotation plus translation [as in the case of Pierre and London], or even ‘disquotation alone’ (or: disquotation plus homophonic translation) [as in the case of Peter and Paderewski]” (p. 269).

As Kripke emphasizes, his aim is not to vindicate Millianism, but merely to argue that – given that the kind of counter-intuitive results in question are already generated by the application of principles the rejection of which would itself be counter-intuitive – “it

would [in the present state of our knowledge] be foolish to draw any conclusion, positive or negative, about substitutivity” (ibid.).

For present purposes, it is not necessary to examine the plausibility of the general line of argument just described. What is more relevant is the fact that, despite Kripke’s motivations for presenting the Paderewski case, this case might be (and also has been) taken to favour what I have called the fine-grained view – the view that the primary objects of knowledge are sensitive to the kind of difference between Norman’s way of thinking about what he calls ‘crimson’ and the way in which people who learned the word ‘crimson’ in the usual way think about this colour – over what I have called the coarse-grained view (which is simply the negation of the former). The reason is that, even though the Paderewski case might seem to present a genuine puzzle both for proponents of the coarse-grained view and for proponents of the fine-grained view, it might still seem that the puzzle, considered from the perspective of proponents of the coarse-grained view, is a more serious one than considered from the perspective of proponents of the fine-grained view. That is, proponents of the fine-grained view may hold that the puzzle merely concerns the relatively superficial issue of the formulations of attitude *reports*, while at the level of the objects of the attitudes, no puzzle arises, since, due to the fine-grainedness of what is believed by Peter when assenting to ‘Paderewski had musical talent’ and ‘Paderewski had no musical talent’, there is no contradiction at the level of the objects of the attitudes. In contrast, for proponents of the coarse-grained view, a contradiction does not only loom at the level of attitude reports, but also at the level of the objects of the attitudes. And this looming contradiction at the more substantial level of the objects of the attitudes has indeed led philosophers to prefer the fine-grained view, even without endorsing any particular solution to the puzzle at the level of attitude reports.

For example, Larson & Ludlow 1993 present a theory of propositions (conceived of as the objects of the attitudes) similar to the one presented in Bealer 1998, which has been mentioned in the previous section as an example of a theory of propositions resulting out of the combination of the fine-grained version of the coarse-grained view and the view that propositions are the primary objects of the attitudes. Like Bealer, Larson & Ludlow conceive of the constituents of propositions as involving referents of expressions as well as the expressions themselves. As already mentioned in the above description of Bealer’s proposal, construing propositions in this way ensures that the resulting theory is one according to which uses of coextensive names like ‘Cicero’ and

‘Tully’, as well as uses of homophonous names with different referents, contribute different entities to the objects of the attitudes. And without further specification, one would expect the intended theory to be one according to which Peter’s use of the name ‘Paderewski’ in his assent to ‘Paderewski had musical talent’ and to ‘Paderewski had no musical talent’ contributes the same entity to the respective proposition. However, Kripke’s scenario has led Larson & Ludlow to treat this issue differently. That is, the scenario has led them to endorse a fine-grained view regarding names. The most self-suggesting way of doing this, without giving up the general framework of their theory, is to treat the use of ‘Paderewski’ in Peter’s assent to ‘Paderewski had musical talent’ and ‘Paderewski had no musical talent’ as the use of two different names. Accordingly,

“we suggest that there are actually two names here, Paderewski_I and Paderewski_{II}, and that the reports in (36) [i.e., ‘Ralph believes [Paderewski is shy]’ and ‘Ralph does not believe [Paderewski is shy]’] are distinguished analogously to those in (35) [i.e., to ‘Max believes [Bill is a flier_I]’, ‘flier_I’ referring to someone who flies, and ‘Max believes [Bill is a flier_{II}]’, ‘flier_{II}’ referring to an advertising circular]” (Larson & Ludlow 1993, p. 319).

It seems quite plausible to assume that Kripke’s case might be modified so as to involve a predicate term, like ‘being crimson’, instead of a proper name. And there seems to be no good reason to treat one kind of term differently from the other in the face of the respective kind of scenario. Accordingly, those who feel themselves forced, in the face of the Paderewski case, to treat Peter’s uses of the name ‘Paderewski’ as contributing different entities to what is believed or known would have equally good reasons to regard Norman’s and the others’ use of the term ‘crimson’ as contributing different entities to what is believed or known.

As described in section 3.2, certain *prima facie* considerations favour the fine-grained version of the coarse-grained view over any version of the fine-grained view. As it has now occurred, considerations regarding the Paderewski case might be taken – and indeed have been, as shown by the example of Larson & Ludlow – as outweighing this kind of *prima facie* considerations. In the remainder of this section, I will argue against this way of reacting to Kripke’s scenario.

As already mentioned above, the reason for proponents of the fine-grained view to regard the Paderewski case as favouring the fine-grained view over the coarse-grained view might be that the puzzle presented is a more serious one considered from the perspective of proponents of the coarse-grained view than considered from the perspective of proponents of the fine-grained view. That is, while from the perspective

of proponents of the fine-grained view, the puzzle merely concerns what might be regarded as the relatively superficial issue of the formulation of attitude *reports*, proponents of the coarse-grained view will have to deal with the threat of a contradiction at the supposedly more substantial level of the attitudes themselves. These considerations, however, rely on the assumption that accepting there to be a contradiction at the level of the primary objects of Peter's beliefs is utterly implausible – an assumption that, as I will now argue, is unwarranted.

As cited above, Kripke's argument against attributing to Peter both the belief that Paderewski had musical talent and the belief that Paderewski had no musical talent rests on the claim that "surely anyone, leading logician or no, is in principle in a position to notice and correct contradictory beliefs if he has them" (Kripke 1979, p. 257). This claim is indeed plausible when understood as a claim about the evaluation of belief reports. That is, surely anyone, leading logician or no, is in principle in a position to notice superficial contradictions of the kind contained in 'Peter believes that Paderewski had musical talent and that Paderewski had no musical talent'. However, taking Kripke's claim literally, as a claim about the detection of contradictions in the objects of beliefs, it is not at all clear why it should be accepted without further support. That is, Kripke's claim, understood literally, seems to simply presuppose that beliefs are cognitively individuated, and thereby (together with certain other assumptions regarding the cognitive accessibility of one's own ways of grasping, thinking, etc.) cognitively accessible, entities. David Lewis makes a very similar point in his discussion of the case of Pierre and London:

"When we characterise the content of belief by assigning propositional (or other) objects, are we characterising an inner, narrowly psychological state of the believer? [...] Or are we characterising partly the believer's inner state, partly the relations of that state to the outer world?^{FN} If it is the latter, [...] Kripke's puzzle vanishes. For if the assignment of propositional objects characterises more than the believer's inner state, then there is no reason to suppose that a leading philosopher and logician would never let contradictory beliefs pass, or that anyone is in principle in a position to notice and correct contradictory beliefs if he has them." (Lewis 1981, pp. 288f.)

Thus, the claim that the Paderewski case provides any relatively serious puzzle for proponents of the coarse-grained view appears unwarranted, as soon as one drops the unsupported assumption that beliefs are cognitively individuated.²² Moreover, while

²² In a section called "Preliminaries: substitutivity", Kripke points out that "*de re* beliefs [...] do not concern us in this paper." (Kripke 1979, p. 242) Lewis considers this announcement as a possible, and charitable, explanation of why Kripke seems to presuppose, in his argument against attributing contradictory beliefs to Peter, that beliefs are cognitively individuated: "I would like to think that this

proponents of the fine-grained view – who claim that the primary objects of Peter’s beliefs do not contain contradictions – still have to explain why it seems correct to attribute both individual beliefs to Peter, proponents of the coarse-grained view are not in need of giving such an explanation, since the correctness of the attribution of the two individual beliefs to Peter is not in conflict with, but rather follows from, their view regarding the primary objects of belief. So at this point, proponents of the coarse-grained view may even seem to be a bit better off in the face of Kripke’s scenario than proponents of the fine-grained view.

However, Lewis’ claim that, given certain assumptions regarding the objects of belief, Kripke’s puzzle vanishes is not quite accurate. As stated above, not simply presupposing that beliefs are cognitively individuated already undermines, without further support for such a view on beliefs, the plausibility of Kripke’s claim that anyone is in principle in a position to notice and correct contradictory beliefs if he or she has them. Not simply presupposing the kind of view on beliefs just mentioned thereby also undermines the reason Kripke provides for not attributing to Peter both the belief that Paderewski had musical talent and the belief that Paderewski had no musical talent. That is, given a view according to which beliefs are not cognitively individuated, it may well seem harmless to attribute to Peter both of these beliefs at the same time. What would, of course, *not* seem harmless is to attribute to Peter both of these beliefs, and, additionally, to infer from the claim that Peter believes that Paderewski had no musical talent that Peter does not believe that Paderewski had musical talent. As Kripke notes, this latter claim, together with the claim that Peter believes that Paderewski had musical talent, would constitute a contradiction which is “not merely in Peter’s judgment, but in our own” (Kripke 1979, p. 258). Kripke calls the principle generating the kind of

[i.e., a narrowly psychological conception of beliefs] was what Kripke intended in instructing us to consider belief *de dicto*.” (Lewis 1981, p. 289) However, explaining Kripke’s refusal to attribute contradictory beliefs to Peter in this way is rather implausible. First, while Kripke uses the term ‘de re’ such that de re contexts are “by definition subject to a substitutivity principle for both names and descriptions [i.e., to a principle according to which coextensive names and descriptions can be substituted in belief reports *salva veritate*]” (Kripke 1979, p. 242), it becomes quite clear in the remainder of the paper that he does *not* use the term ‘de dicto’ such that de dicto contexts are by definition subject to a *non*-substitutivity principle for coextensive names and descriptions. Instead, he uses the term such that de dicto contexts are simply *not* by definition subject to the kind of substitutivity principle in question. Using the term in the first of the two ways just mentioned would amount to stipulating that the only contexts to be considered in the paper are ones in which Millianism is false – a stipulation obviously not intended by Kripke. Second, even if Kripke made such a stipulation, the non-substitutivity of coextensive terms would still leave open whether beliefs are cognitively individuated in the relevant sense, that is, in the sense necessary for the *single* name (and not just two different, coextensive names) ‘Paderewski’ to contribute different entities to the objects of belief. For these reasons, I will not deal with the de re/de dicto distinction any further.

inference in question the ‘strengthened disquotational principle’. But not only does he not explicitly endorse this principle, he also provides no reasons for accepting it, and the mere possibility of applying it, without any reasons for accepting it, does clearly not threaten the attribution of the beliefs that Paderewski had musical talent and that Paderewski had no musical talent to Peter. The reason why Lewis is wrong in claiming that, given certain assumptions regarding the objects of belief, Kripke’s puzzle vanishes is, instead, the fact that it seems incorrect to attribute to Peter the belief that Paderewski had both musical talent and no musical talent. And an explanation would be needed why an inference from the claim that Peter believes that Paderewski had musical talent and the claim that Peter believes that Paderewski had no musical talent to the counter-intuitive attribution just mentioned is impermissible. Without such an explanation, the puzzle has clearly not vanished.

At this point, proponents of the fine-grained view might claim that the Paderewski case favours, in the end, their view over the fine-grained version of the coarse-grained view, since the fine-grained view is compatible with a self-suggesting way of resolving the remainder of the puzzle described in the previous paragraph – a way not available to proponents of any version of the coarse-grained view. The self-suggesting way I have in mind here consists in endorsing a contextualist view according to which attributions like ‘Peter believes that Paderewski had musical talent’ and ‘Peter believes that Paderewski had no musical talent’ both often express truths, not because Peter holds contradictory beliefs in any substantial sense, but merely because the term ‘Paderewski’ contributes different entities to what is attributed to Peter, depending on the context of utterance. Accordingly, the contextualist view in question could account for the incorrectness of attributions like ‘Peter believes that Paderewski both had musical talent and no musical talent’, since – so proponents of the view would hold – none of the two different entities contributed by the term ‘Paderewski’ to what is attributed to Peter is such that Peter believes of it that it both had musical talent and had no musical talent. The view in question might thereby be regarded as providing an elegant solution to the remainder of the puzzle – a solution obviously not available to proponents of any version of the coarse-grained view, since it involves, as an essential part of it, the claim that Peter’s uses of the name ‘Paderewski’, when assenting to ‘Paderewski had musical talent’ and ‘Paderewski had no musical talent’, contribute the same entity to what is believed.

The contextualist view in question should, most plausibly, be regarded as involving an analogous claim regarding Norman’s and the others’ uses of the term ‘crimson’. It

might thus seem that, in the end, considerations regarding Kripke's scenario outweigh the prima facie considerations presented in section 3.2, and that they thereby vindicate the argument from fineness of grain. However, this conclusion would be too hasty. For under closer inspection, the seemingly elegant solution to the remainder of the puzzle just described is far less plausible than it might first appear. Consider two persons, both sincerely assenting to 'Paderewski had musical talent'. One of them only knows Paderewski as a musician, while the other only knows him as a statesman, nonetheless, for whatever reasons, sincerely assenting to the sentence in question. Endorsing the contextualist proposal described above would commit one to claiming that the attribution 'Both persons believe that Paderewski had musical talent' expresses a falsehood, since, according to the proposal in question, whatever fine-grained entity the term 'Paderewski' contributes to what is attributed via the above sentence, it is, due to the two persons' different kinds of familiarity with Paderewski, not an entity that figures in *both* persons' respective beliefs. But clearly, being committed to claiming that 'Both persons believe that Paderewski had musical talent' expresses a falsehood should be regarded as an unwelcome consequence of endorsing the contextualist view in question. Moreover, if proponents of the fine-grained view wish to deal with Kripke's puzzle by accepting this consequence, then proponents of a version of the coarse-grained view might, in response, propose to deal with the puzzle by accepting that the attribution 'Peter believes that Paderewski both had musical talent and had no musical talent' expresses a truth. Regarding degrees of counter-intuitiveness, both strategies seem to be more or less on a par.

My considerations regarding the alleged threat posed to the coarse-grained view by the Paderewski case can be summarized as follows. First, it has turned out that the worry according to which Kripke's puzzle is a more serious one when considered from the perspective of proponents of the coarse-grained view appears far more harmless than Kripke's own considerations suggest. That is, as soon as one drops the assumption that beliefs are cognitively individuated and cognitively accessible, being committed to claiming that there is a contradiction in the objects of Peter's beliefs appears far more harmless than described by Kripke. It has, however, also turned out that dealing with the puzzle in this way commits one, at least prima facie, to accepting the rather unwelcome consequence that attributions like 'Peter believes that Paderewski had both musical talent and no musical talent' express truths. I have considered a self-suggesting way for proponents of the fine-grained view – namely, a certain contextualist version of the

view – to avoid being committed to this claim. But as pointed out, dealing with the puzzle in this way would lead one to accept certain consequences that should be regarded as similarly unwelcome. Thus, at least as far as my above considerations go, Kripke’s case remains puzzling, both from the perspective of proponents of the coarse-grained view and from the perspective of proponents of the fine-grained view. Since the consequences that proponents of the two views are committed to when dealing with the puzzle do not significantly differ in degrees of absurdity, the puzzle does not seem to favour the fine-grained view over the coarse-grained view. My considerations in this section do certainly not disprove that there is some way available to proponents of the fine-grained view, but not to proponents of the coarse-grained view, to elegantly avoid accepting the unwelcome consequences of their respective ways of dealing with the puzzle. But I take it that the burden of proof regarding such a claim lies with proponents of the fine-grained view. As long as no support has been provided for this claim, there is good reason to assume that the puzzle does not constitute a case for the fine-grained view that could outweigh the *prima facie* considerations in favour of the fine-grained version of the coarse-grained view presented in section 3.2.²³

3.5. Conclusions regarding the truth of P1

In section 3.2, I have described how certain intuitions regarding sameness of what is known constitute a *prima facie* case for what I have called the fine-grained version of the coarse-grained view – a position according to which expression pairs like ‘Hesperus’ and ‘Phosphorus’ contribute different entities to the objects of the attitudes, while the term ‘crimson’, as used by Norman and those who learned the term in the usual way, contributes the *same* entity to the objects of the attitudes. In the preceding two sections, I have argued that certain worries regarding the fine-grained version of the

²³ Powell 2012 uses the notion of a partially defined predicate, developed in Soames 1999 (see esp. pp. 163ff.), in order to present a solution to the case of Pierre and London according to which the truth of each of the problematic belief attributions in question should not be denied, but rejected – where rejecting differs from denying in that rejecting the truth of an attribution does not commit one to accepting that the attribution is not true. Applying the proposal in question to the Paderewski case would result in rejecting both the claim that Peter believes that Paderewski had musical talent and the claim that Peter believes that Paderewski had no musical talent, thereby neither accepting that these claims are true nor that they are not true. I will not examine the prospects of Powell’s proposal. Instead, it suffices, for present purposes, to note that the proposal, if promising at all, would obviously be equally available to proponents of the fine-grained view and to proponents of the fine-grained version of the coarse-grained view in order to avoid the respective unwelcome commitments described above.

coarse-grained view do not outweigh the considerations presented in section 3.2. Showing that the two worries considered do not outweigh the prima facie considerations in favour of the fine-grained version of the coarse-grained view does of course not amount to showing that no other considerations would, in the end, favour the fine-grained view over the fine-grained version of the coarse-grained view. However, I take it that the burden of proof now lies with proponents of the fine-grained view and that, given the considerations in the preceding three sections, there is, as long as no further arguments are presented, good reason to suppose that the fine-grained view is false, and that the argument from fineness of grain, presented in section 3.1 as an argument against P1, is thereby unsound.

In the face of the above considerations, opponents of P1 might try to resort to a view according to which the fineness of grain of the objects of the attitudes varies across contexts. That is, they might try to resort to a view according to which in some contexts of utterance, reports about what someone believes or knows attribute to this person an attitude of a degree of fineness of grain corresponding to the fine-grained view, while in other contexts of utterance, attitudes of a lower degree of fineness of grain are attributed, even though the formulations used do not differ. (This kind of contextualist position is endorsed by Bealer 1998, p. 9.) Accordingly, opponents of P1 might resort to the claim that the argument from fineness of grain is sound at least from the perspective of certain contexts. From the perspective of these contexts, one would, for instance, express a falsehood by uttering, regarding someone knowing Paderewski only as a Polish pianist and someone knowing Paderewski only as a Polish statesman, ‘Both know that Paderewski was Polish’ – since whatever fine-grained entity the name ‘Paderewski’ contributes to what is attributed in the knowledge report, it is not an entity that figures in the objects of knowledge of both of these persons.

I grant that if the contextualist position just described is correct, then the argument from fineness of grain might well be successful from the perspective of certain contexts, and accordingly, the argument from analogy might only be successful from the perspective of other contexts. However, not only is the kind of contextualism in question at best a minority view, there has also not yet been provided any attempt at describing the possible mechanisms behind the postulated kind of context shifts. In the debate about contextualism regarding knowledge, the kind of description in question usually concerns, more or less directly, the salience of certain possibilities of error. As long as no analogous explanation has been made plausible regarding the kind of

contextualism described in the previous paragraph, it may simply be assumed that the unsoundness of the argument from fineness of grain is not a matter of context.²⁴

It may be pointed out that defending P1 against one particular counter-argument does not yet amount to a positive case in favour of it. However, my above defense of P1 against the argument from fineness of grain involves a positive case for a certain view regarding the objects of the attitudes – namely, for the fine-grained version of the coarse-grained view, according to which the objects of the attitudes are not fine-grained enough to be sensitive to the cognitive differences between Norman and the others regarding the term ‘crimson’. And I claim that a positive case for this view about the objects of the attitudes also provides a good deal of positive support for P1.

Note that one can easily construe the example of Mary and Norman such that Norman is not only *as* reliable in distinguishing crimson from non-crimson things as people who learned the term ‘crimson’ in the usual way and do not know anything about the relevant microphysical surface properties, but even *more* reliable in this respect than most people of the latter kind (given, of course, that Norman is equipped with a microscope²⁵). Norman’s way of distinguishing crimson from non-crimson things is, for example, not vulnerable to a proper adaptation of the eyes to the relevant lighting conditions. And one may construe the example such that Norman is maximally competent regarding his specific method of distinguishing crimson from non-crimson things – as competent as any human being could get in this respect. (One may also construe the example such that Norman is not only maximally competent in using properly functioning microscopes for the purpose in question, but also in judging whether a given microscope is properly functioning.) And recall that he is aware of the fact that he is not able to distinguish crimson from non-crimson things on the basis of looking at them with the unaided eye, so that this deficiency does not lead him to make false judgments about whether an object is crimson.

²⁴ A possible candidate for a description of the mechanisms behind the kind of context-shift in question would be one according to which the shifts are triggered by significant changes in the salience of the general possibility of failing to realize that a supposedly homonymous name or predicate in fact only refers to one single individual or property; or maybe already by changes in the salience of the general possibility of thinking of a given referent in different ways. As already stated above, I take it that the burden of proof lies with those who claim that candidates like the ones just mentioned can figure in a plausible explanation of why, from the perspective of some contexts, the fine-grained view is correct.

²⁵ One may even conceive of a version of the case in which Norman can, due to some strange kind of genetic mutation, focus his eyes in such a way as to enable him to detect microscopic surface properties without using a microscope. Such a scenario may be described as one in which he is able to distinguish crimson from non-crimson things on the basis of their microscopic surface properties with the unaided eye.

As pointed out in section 3.1, P1, if it is to be relevant regarding debates on a priori truth, should be interpreted as the claim that in virtue of Norman's competence just described, he understands or grasps the entity contributed by the word 'crimson' to the object of truth corresponding to the sentence 'All crimson things are red' – whatever kind of entity the primary objects of truth turn out to be. Now, given that,

- (i) as argued for in the preceding three sections, the entity contributed to the primary objects of truth by the word 'crimson' are not fine-grained enough for their individuation to be sensitive to the cognitive differences between Norman and the others regarding this term, and given that,
- (ii) as described above, Norman is at least as reliable as the others in distinguishing crimson from non-crimson things, and finally, given that
- (iii) the fact that Norman and the others use the same word in order to refer to the colour crimson is not the coincidental product of some deviant chain of reference (but rather the result of a transmission of the word from people who can only detect crimson things in the usual way to Mary, who can detect crimson things both in the usual and in the microscope-involving way, to Norman, who can only detect crimson things in the microscope-involving way),

it is not at all clear why Norman should be regarded as having less understanding or grasp of the entity contributed by the word 'crimson' to the object of truth corresponding to the sentence 'All crimson things are red' than those who are only able to detect crimson things in the usual way. It might, of course, be claimed that Norman has a poorer understanding or grasp of the entity in question than Mary, who is able to detect crimson things both in the usual and the microscope-involving way. But in the face of the above considerations, attributing to those who can only detect crimson things in the usual way understanding or grasp simpliciter of the entity in question commits one to the same attribution regarding Norman. Both Norman and the others have a method of reliably distinguishing crimson from non-crimson things; and given the above considerations regarding the primary objects of truth, none of these methods should be regarded as privileged over the other when it comes to attributing understanding or grasp of the entity in question. Since proponents of the apriority of what is expressed by 'All crimson things are red' would probably in general be

unwilling to deny people who can only detect crimson things in the usual way understanding or grasp simpliciter of the entity in question, they should thereby also be unwilling to make such a denial regarding Norman.

Thus, my considerations in favour of the fine-grained version of the coarse-grained view do not only provide a case for the unsoundness of the argument from fineness of grain, they also provide a positive case for P1, if the latter is interpreted in the way relevant concerning debates about apriority. As mentioned above, the burden of proof regarding the truth or falsity of the fine-grained version of the coarse-grained view lies with those who deny the truth of this view. Accordingly, the burden of proof regarding the truth or falsity of P1 now lies with those who deny that P1 is true.

3.6. Summary and outlook

In the preceding sections, I have defended P1 – or, more exactly, a reading of P1 relevant regarding debates about apriority – against a certain objection, namely, against the argument from fineness of grain, and I have, in the course of this defense, also provided a positive case for the relevant reading of P1. In section 3.1, I have developed the argument from fineness of grain, and I have explained in how far this argument, if successful, would undermine P1. In section 3.2, I have introduced what I have called the fine-grained version of the coarse-grained view – a view whose truth would undermine the soundness of the argument from fineness of grain, and according to which Norman’s use of the word ‘crimson’ contributes the same entity to the objects of the attitudes as the use of this word by those people who learned it in the usual way, but according to which the use of different words still contributes different entities to the objects of truth – and I have presented some prima facie considerations favouring the view over its competitors. In sections 3.3 and 3.4, I have argued that two worries concerning the view – worries connected to the practice of translation, and worries connected to Kripke’s Paderewski case – should not be regarded as outweighing the prima facie considerations presented in section 3.2. I have thereby shifted the burden of proof to those who claim that the fine-grained version of the coarse-grained view is false. Finally, in section 3.5, I have argued that my considerations in favour of the fine-grained version of the coarse-grained view do not only constitute a case against the soundness of the argument from fineness of grain, but also a positive case for P1, if the latter is interpreted in a way

relevant regarding debates about apriority. I have thereby shifted the burden of proof to those who deny that P1 is true.

Having defended the argument to C1 in chapters 2 and 3, I will now, in the following chapter, provide a defense of the remaining premise of the argument from analogy.

4. Examining the argument from analogy, part III: defending P3

In the previous two chapters, I have argued in favour of the first half of the argument from analogy, that is, in favour of the argument to C1. Since C1 and P3 clearly entail C2, P3 is the only part of the argument from analogy that remains to be defended. Accordingly, my aim in the present chapter is to provide considerations in favour of P3. In section 4.1, I will examine several candidates for a general principle instantiated by – and thereby entailing – P3. If any such principle turns out to remain unrefuted in the face of possible counterexamples, this will constitute a strong case for P3. However, instead of endorsing any specific version of such a principle, I will argue that the only remaining possible counterexamples are those involving a significant concession by opponents of P3 – namely, the concession that it is in principle possible for an experience-token to play both an enabling and an evidential role. In section 4.2, I will present an example that could – independently of the observations provided in section 4.1 – very plausibly be counted as involving experience that plays this kind of double-role. Moreover, I will argue that there is no promising candidate for a distinguishing criterion in virtue of which the relevant experience in the example in question, but not also some experience in the original version of the case of Norman, plays a double-role. Furthermore, I will argue in favour of a criterion in virtue of which not only the relevant experience in the example in question, but also some of the experience in the original version of the case of Norman plays a double-role. Finally, in section 4.3, I will provide an additional argument in favour of P3, supporting the considerations provided in the previous sections.

4.1. Searching for the principle behind P3

As already mentioned in the preliminary presentation of the argument from analogy in section 1.4, the initial motivation for assuming P3²⁶ stems from the fact that Norman's

²⁶ By way of reminder: P3 is the claim that if Norman's experiences regarding the superficial looks of crimson things play an evidential role in the modified version of the scenario (the one involving Mary) in his knowledge that all crimson things are red, then these experiences also play an evidential role in the *original* version of the scenario in his knowledge that all crimson things are red.

knowledge of R1 – that is, his knowledge that all crimson things are red – in the modified version of the scenario (the version in which he learns to individuate crimson things via the information provided by Mary) and his knowledge of R1 in the original version of the scenario (the version in which he comes to understand the term ‘crimson’ in the ordinary way) are similar in certain respects. That is, in both versions of the scenario, the knowledge in question is acquired by Norman’s exercise of his reliable offline skill of representing to himself the superficial colour properties of crimson things and red things. And in both versions of the scenario, this reliable offline skill partly stems from Norman’s visual experiences regarding the superficial looks of crimson things.

As already conceded in section 1.4, stating these similarities does of course not yet constitute a convincing case for P3. One possible way of providing a convincing case for P3 would, for instance, be to identify the latter as an instance of a more general principle regarding similarities of roles of experience in certain kinds of cases, and to show that this principle resists possible counterexamples. In the present section, I will examine a number of candidates for a principle which might play this role.

A first self-suggesting candidate is the following one:

(P3_{g1}) If some kind of experience plays an evidential role in someone’s knowledge of p, then this kind of experience also plays an evidential role in any other instance of knowledge of p whose acquisition is at least partly based on this kind of experience.

The expression ‘based on’ used in P3_{g1} is supposed to be understood in a wide sense, so that if someone’s knowledge of p is at least partly based on her understanding of a certain word, and if a certain kind of experience plays the role of providing for, or of significantly contributing to, the subject’s understanding of this word, then the subject’s knowledge of p is at least partly based on the kind of experience in question, in virtue of the circumstances just described. Given this wide understanding of the expression ‘based on’, not only Norman’s knowledge of R1 in the modified version of the scenario, but also his knowledge of R1 in the original version of the scenario is at least partly based on experiences regarding the superficial looks of crimson things. P3_{g1} is therefore instantiated by P3. Thus, if P3_{g1} is true, P3 is true as well. However, P3_{g1} should not be

expected to play the desired role in a defense of P3, because it turns out to be false in the light of the following example²⁷ (case 1 in appendix B²⁸).

Consider Susan, who is able to distinguish crimson from non-crimson things as reliably as most other people on the basis of looking at them with the unaided eye, and who therefore counts as understanding the term ‘crimson’. She came to understand the term in the course of looking at various colour samples during a school lesson, and being told by the teacher which of them count as crimson (and which of them should be regarded as borderline cases). Later, she came across a newspaper report about the 1986 football team of Washington State University attached by (printings of) photos of the team members wearing their home jerseys. Given her good visual memory of what she learned in the school lesson just described, Susan comes to know – by looking at the photos, which authentically depict the jerseys’ colour, and which are accompanied by a caption stating that the players are wearing their home jerseys on them – that the home jerseys of the 1986 football team of Washington State University were crimson. Since the photos are not accompanied by any description of the jerseys’ colour, looking at the photos does not contribute to Susan’s understanding of the term ‘crimson’. Rather, the visual experiences she has while looking at the photos play a clearly evidential role in her knowledge that the home jerseys of the 1986 football team of Washington State University were crimson.

Next, consider Anna, who, like Susan, is able to distinguish crimson from non-crimson things as reliably as most other people on the basis of looking at them with the unaided eye, and who therefore counts as understanding the term ‘crimson’. She first became familiar with the term at some particular time in her childhood, while looking at (printings of) the same photos of the 1986 football team of Washington State University already mentioned above. At this occasion, she was (truthfully) told that the colour of the jerseys is one of the paradigm samples of crimson. Since Anna has not been given many more ostensive definitions of the term ‘crimson’ afterwards, looking at these photos significantly contributed to her understanding of the term. Over the years, however, she completely forgot about the occasion and the fact that her understanding of the term ‘crimson’ (which persists over the years) stems to a significant extent from

²⁷ I am indebted to John Horden for presenting me with a similar kind of example.

²⁸ In the course of the present and the following section, I will construe a series of more or less similar examples, all of which involve the same main protagonists, as well as the same piece of knowledge acquired by these protagonists. I list these examples in tabular form at the end of the dissertation, in appendix B, in order to make retrospective references to them more reader-friendly.

seeing the photos in question. Moreover, she did not come to know at this occasion that the home jerseys of the 1986 football team of Washington State University were crimson, for she merely paid attention to the photos themselves and did not read the article or the caption. Years later, she *did* come to know that the home jerseys of the 1986 football team of Washington State University were crimson – not by seeing photos, though, but instead via relying on oral testimony. Clearly, the visual experiences Anna had while looking at the photos merely contributed to Anna’s coming to understand the term ‘crimson’, and play therefore no evidential role in her testimony-based knowledge. Claiming otherwise, without any surprising argument, would be completely ad hoc.

Unfortunately, applying $P3_{g1}$ to case 1 (‘case 1’ referring to the pair of instances of knowledge just described) yields the unwelcome result that, given the hard-to-deny fact that the experience of seeing the photos plays an evidential role in Susan’s knowledge that the home jerseys of the 1986 football team of Washington State University were crimson, it also plays an evidential role in Anna’s knowledge of this proposition (or whatever is considered as the primary object of knowledge). $P3_{g1}$ applies in this way since, given the intended wide reading of the expression ‘based on’ described above, not only Susan’s knowledge, but also Anna’s knowledge counts as being based on the visual experience of looking at the photos. Moreover, the visual experience Anna had while looking at the photos should count as the same *kind* of experience as the visual experience Susan had while looking at the photos. Any significant narrowing of the reading of the term ‘kind’ in the formulation of $P3_{g1}$, to the effect that $P3_{g1}$ would no longer apply to the two cases in the way described above, would not only appear ad hoc, but would also have the unwelcome effect that $P3_{g1}$ would no longer apply in the desired way to the two cases involving Norman (– since the relevant experiences of seeing crimson things with the unaided eye in the two cases involving Norman are described in much more general terms than the relevant experiences in the cases of Susan and Anna; moreover, the cases of Susan and Anna may simply be modified such that their visual experiences of seeing the photos are even more similar, e.g. regarding duration, distance, etc., and they may even be modified such that the same exemplar of the newspaper figures in the two cases.) Thus, $P3_{g1}$ is false and can therefore not figure in a case for P3.

As just mentioned, any promising attempt at modifying $P3_{g1}$ such that it does no longer classify the relevant visual experience in Anna’s knowledge as evidential has to

be such that it nevertheless still classifies the relevant visual experience in the original version of the case of Norman as evidential. For such a modification to be found, one might look at differences between the cases of Susan and Anna that do not occur between the two cases involving Norman. A modified version of $P3_{gl}$ sensitive to such kind of difference may then no longer yield the undesired analogy-claim regarding the cases of Susan and Anna, while still yielding the desired analogy-claim regarding the two cases involving Norman.

Obviously, the fact that the cases of Susan and Anna involve, as their main protagonists, different persons – Susan and Anna – while the two cases involving Norman involve the same main protagonist, does not constitute a kind of difference that could be a basis for a promising modification of $P3_{gl}$. For it has been an arbitrary choice to reuse the name ‘Norman’ in the description of the modified version of the case. Whether the name ‘Norman’ refers, strictly speaking, to the same person in the two cases is, of course, a controversial issue (given that the two cases happen in different possible worlds). And regardless of whether the person called ‘Norman’ in the description of the modified version of the case and the main protagonist in the original version of the case are identical, or merely what David Lewis calls ‘counterparts’²⁹, the respective difference between the two pairs of cases can simply be overcome by replacing Anna with someone who, analogously to the other pair of cases, is either identical to or a counterpart of Susan.

Instead, another difference between the two pairs of cases appears more promising when it comes to attempting to find a basis for a useful modification of $P3_{gl}$. This difference is constituted by the fact that, while the two cases involving Norman might be construed such that Norman’s acquisition of the relevant piece of knowledge in the original version of the case does not involve any kind of experience that does not also figure in the modified version of this case (that is, the two cases might be construed such that Norman comes to know about the superficial looks of crimson and red things by seeing the same kinds of objects under the same kinds of conditions, etc.), the acquisition of the relevant piece of knowledge in the case of Anna is based on a certain kind of experience – namely, a testimonial one – which does not figure in Susan’s

²⁹ See Lewis 1968 for a general introduction of the counterpart relation, and Lewis 1971 for an application of the framework to puzzles about personal identity. For matters of simplicity, and since it is of no importance for present purposes, I have described, and will continue to describe, the two cases, as well as others, in a way presupposing cross-world identities.

acquisition of the relevant piece of knowledge. Thus, the following modification of P3_{G1} appears worth testing:

(P3_{g2}) If some kind of experience plays an evidential role in someone's knowledge of p, then this kind of experience also plays an evidential role in any other instance of knowledge of p *which is not at least partly based on any kind of experience that does not also figure in the first instance of knowledge of p* and whose acquisition is at least partly based on this kind of experience.

As already mentioned, this version of the principle applies in the desired way to the pair of cases involving Norman, since the two cases might be construed such that no kind of experience figuring in the original version of the case does not also figure in the modified version. The relevant question now is whether the pair of cases involving Susan and Anna can be modified accordingly, preserving the unwelcome application of the principle described above.

And indeed, P3_{g2} turns out to be false in the light of the following modification of case 1 (case 2 in appendix B). Imagine that, instead of coming to understand the term 'crimson' in the course of a school lesson, Susan comes to understand the term in a way involving the same kind of testimonial experience that plays an evidential role in Anna's knowledge that the home jerseys of the 1986 football team of Washington State University were crimson. That is, imagine, for example, that Susan comes to understand the term 'crimson' by being told that it is the colour of the home jerseys of the 1986 football team of Washington State University, and by, say, subsequently looking at an original exemplar of one of the jerseys in a football museum. Imagine further that, over the years, Susan completely forgets about the way in which she came to understand the term, and also about the fact that the term refers to the colour of the home jerseys of the 1986 football team of Washington State University. As in the original version of the case, she later (re)acquires the relevant piece of knowledge by reading the newspaper article and looking at the photos attached to it. As in case 1, the visual experiences she has while looking at the photos play a clearly evidential role in her acquisition of the relevant piece of knowledge. Moreover, all kinds of experience figuring (one way or another) in Anna's knowledge in case 2 also figure in Susan's knowledge.³⁰ And since,

³⁰ This latter claim is actually not quite accurate, since the testimonial experience of being told that the colour of the jerseys on the photos is one of the paradigm samples of crimson, which figures in Anna's

as already stated above, the visual experiences Anna had while looking at the photos play no evidential role in her testimony-based knowledge, applying $P3_{g2}$ to case 2 yields the same unwelcome result already yielded by $P3_{g1}$ applied to case 1.

At this stage of developing a formulation of a principle behind $P3$, it is necessary to be particularly cautious. Having introduced case 2, in which Susan's experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson significantly contributes to her coming to understand the term 'crimson', one might be tempted to regard this kind of experience as playing an evidential role in Anna's knowledge, but a merely *enabling* role in Susan's knowledge. And one might be tempted to use this assumption as a basis for the following reformulation of $P3_{g2}$:

($P3_{g3}$) If some kind of experience plays an evidential role in someone's knowledge of p , then this kind of experience also plays an evidential role in any other instance of knowledge of p which is not at least partly based on any kind of *evidential experience that does not also play an evidential role* in the first instance of knowledge of p and whose acquisition is at least partly based on this kind of experience.

This version of the principle still applies in the desired way to the pair of cases involving Norman, for, as already mentioned, the two cases might be construed such that no kind of experience figuring in the original version of the case does not also figure in the modified version, and *if any* of the experience figuring in the original case of Norman plays an evidential role, then the same kind of experience also plays an evidential role in the modified version of the case, given that the two cases are construed similarly enough (– the controversial issue is not whether the experience of seeing crimson things with the unaided eye plays an evidential role in the modified version of the case, but merely whether it plays an evidential role in the original version).

coming to understand the term 'crimson' in case 2, does not figure in Susan's knowledge. Thus, given a sufficiently narrow reading of the term 'kind of experience' in $P3_{g2}$, case 2 is not a counterexample to this principle. However, I simply grant, in favour of opponents of $P3$, that case 2 could somehow be modified such that the kind of testimonial experience just mentioned also appears, as a supplementary factor, in Susan's coming to understand the term 'crimson'.

However, introducing P3_{g3} can only avoid yielding the unwelcome result described above in the discussion of P3_{g1} and P3_{g2} if the experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson, which clearly plays an evidential role for Anna, does not also play an evidential role for Susan. And even though this latter claim may appear *prima facie* plausible, endorsing it would be a grave methodological mistake by defenders of the argument from analogy. This is because the argument from analogy is supposed to constitute a case for the surprising claim that the experience of seeing crimson things with the unaided eye plays an evidential role in the original version of the case of Norman. More exactly, as described in section 1.3, the position endorsed by defenders of the argument for analogy is that the kind of experience in question plays a *double*-role in the original version of the case of Norman, that is, that it plays both an enabling and an evidential role. Putting it in a condensed way, the feature in virtue of which, according to defenders of the argument from analogy, the kind of experience in question plays this double-role consists in the fact that Norman's knowledge that all crimson things are red in the original version of the case is not simply partly enabled by his understanding of the term 'crimson', but it also depends on the specific way in which he came to understand the term. That is, if he had come to understand the term not in terms of the superficial looks of crimson things, but instead in terms of their microphysical surface properties, then, given the way in which he came to understand the term 'red', he would not have been in a position to acquire the knowledge that all crimson things are red via offline reasoning. And this very feature occurs as well in Susan's knowledge in case 2. That is, if Susan had come to understand the term 'crimson' in terms of their microphysical surface properties, she would not have been in a position to acquire the knowledge that the colour of the home jerseys of the 1986 football team of Washington State University is crimson via looking at the photos and reading the caption. Therefore, defenders of the argument from analogy are committed to claiming that the experiences of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson does not only play an evidential role in Anna's knowledge, but also in Susan's knowledge in case 2.³¹ (Note that all this is compatible with my above

³¹ Matters are actually more complicated in light of the fact that some experience of being told, in some way or other, that certain things are crimson would probably also have to be involved in coming to understand the term 'crimson' in terms of the microphysical surface structure of crimson things. But this fact would at best make life easier for defenders of the argument from analogy; for, if the relevant kind of testimonial experience does not play an evidential role in Susan's knowledge in the modified version of

claim that the experiences Anna has while looking at the photos attached to the newspaper article do *not* play an evidential role in cases 1 and 2; this is because, even if Anna had come to understand the term ‘crimson’ in terms of the microphysical surface structure of crimson things, she would still have been in a position to acquire the knowledge that the colour of the home jerseys of the 1986 football team of Washington State University is crimson via relying on testimony.)

The plausibility of regarding the feature just described as a criterion for evidentiality will be discussed in section 4.2 below. For the moment, it is merely important to note that defenders of the argument from analogy do not have at their disposal the claim that the testimonial experience that plays an evidential role for Anna does not also play an evidential role for Susan in case 2. And since, if the testimonial evidence in question plays an evidential role for Susan, P3_{g3} applies in the unwelcome way described above to case 2 (implausibly classifying the visual experiences Anna had while looking at the photos attached to the newspaper report as playing an evidential role in her testimony-based knowledge), defenders of the argument from analogy have to look for a different version of the principle.

Fortunately, the difficulty just described can be avoided relatively easily, by modifying P3_{g3} as follows:

(P3_{g4}) If some kind of experience plays an evidential role in someone’s knowledge of p, then this kind of experience also plays an evidential role in any other instance of knowledge of p which is not at least partly based on any kind of *exclusively*³² evidential (that is, evidential but non-enabling) experience that does not also play an *exclusively* evidential role in the first instance of knowledge of p and whose acquisition is at least partly based on this kind of experience.

the case, P3_{g3} does not yield the unwelcome result described above when applied to case 2. Therefore, I will, for the sake of brevity, leave this point aside here.

³² The difficulty described above could also have been avoided by using the more restrictive formulation ‘...which is not at least partly based on any kind of evidential experience that does not play an *exclusively* evidential role in...’. However, defenders of the argument from analogy are not in a position to claim that the resulting version of the principle would still apply in the desired way to the pair of cases involving Norman. The reason is that those who claim that the experience providing Norman with an understanding of the term ‘crimson’ plays a double-role, and thereby also an evidential role, in the original version of the case are committed to accepting an analogous claim regarding the experience providing Norman with an understanding of the term ‘red’ in the original version of the case. But the fact that the latter kind of experience clearly plays no *exclusively* evidential role in the modified version of the case prevents the version of the principle just considered from applying in the desired way to the two cases.

Like the previous versions of the principle, $P3_{g4}$ applies in the desired way to the pair of cases involving Norman. For, as already mentioned, *if* Norman's knowledge in the original version of the case involves any kind of evidential experience, then the same kind of experience also plays an evidential role in the modified version of the case, given that the two cases are construed similarly enough (– again, the controversial issue is not whether the experience of seeing crimson things with the unaided eye plays an evidential role in the modified version of the case, but merely whether it plays an evidential role in the original version). And if Norman's knowledge in the original version of the case does not only involve some kind of evidential experience, but also some *exclusively* evidential experience – a claim not even defenders of the argument from analogy are committed to –, then the same kind of experience also plays an exclusively evidential role in the modified version of the case, given that the two cases are construed similarly enough.

Moreover, $P3_{g4}$ does not apply in the unwelcome way described above to case 2. For the experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson does not only, as mentioned above, play an evidential role in Anna's knowledge, but it is also quite plausible to assume that this role is an *exclusively* evidential one, since Anna already understands the term 'crimson' before having the testimonial experience in question, as well as – so one may stipulate – good offline reasoning skills corresponding to her way of understanding the term. And if the kind of testimonial experience in question plays an evidential role in case 2 for Susan, then this role will not be an *exclusively* evidential one, since, as stated above, it plays a clearly enabling role in this case. Thus, there is some kind of experience playing an exclusively evidential role in case 2 for Anna, but not for Susan, which keeps $P3_{g4}$ from yielding the unwelcome claim that the experience of looking at the photos attached to the newspaper report, which plays an evidential role for Susan in both cases considered so far, also plays an evidential role for Anna in one of the cases.

Unfortunately, however, $P3_{g4}$ turns out false, in the light of another modification of the case of Susan (case 3 in appendix B). Imagine that, as in case 1, Susan comes to understand the term 'crimson' during a school lesson, by looking at various colour samples and being told by the teacher which of them count as crimson. Imagine further that, in contrast to cases 1 and 2, Susan does not acquire the relevant piece of knowledge merely via looking at the photos and reading the caption, but by also relying on an additional source of evidence, constituted by the same kind of testimony that

figures in the case of Anna. The fact that Susan, in this case, relies on two different kinds of experience is clearly compatible with the claim that both kinds of experience play an evidential role in this case. And since, like Anna, Susan already understands the term ‘crimson’ and possesses good corresponding offline reasoning skills before having the relevant testimonial experience, it is quite plausible to hold that this kind of experience plays an *exclusively* evidential role, not only for Anna, but also for Susan in this case. Thus, given that the kind of testimonial experience in question is the only exclusively evidential experience for Anna, nothing in case 3 keeps P3_{g4} from yielding the unwelcome claim that the experience of looking at the photos attached to the newspaper report, which plays an evidential role for Susan, also plays an evidential role for Anna.

In order to propose a version of the principle that avoids this kind of counterexample, I have to introduce the notion of an evidential burden. This notion shall be defined such that the evidential burden carried by some particular kind of experience in some particular instance of knowledge is determined by the extent to which the piece of knowledge in question is based on some other kinds of evidential experience. Given this definition, the experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson does not carry the same evidential burden for Susan and Anna in case 3, since regarding Anna, the relevant piece of knowledge is not based on any other kind of evidential experience, while regarding Susan, the relevant piece of knowledge is also based, as in cases 1 and 2, on the evidential experience of looking at the photos attached to the newspaper report. Thus, the pair of cases in question does not constitute a counterexample to the following version of the principle:

(P3_{g5}) If some kind of experience plays an evidential role in someone’s knowledge of p, then this kind of experience also plays an evidential role in any other instance of knowledge of p which is not at least partly based on any kind of exclusively evidential experience that does not also play an exclusively evidential role, *or does not carry the same evidential burden*, in the first instance of knowledge of p and whose acquisition is at least partly based on this kind of experience.

Not only does this version of the principle, as explained above, not apply in the unwelcome way described above to case 3, it also still applies in the desired way to the pair of cases involving Norman. For the modification just made could only keep the principle from applying in the desired way to the pair of cases involving Norman if some kind of experience played an exclusively evidential role in the original version of the case – a claim that defenders of the apriority of R1 should of course be unwilling to accept, and which, furthermore, lacks any apparent support.

In the remainder of this section, I will argue that P3_{g5} does not only avoid the kinds of counterexample considered so far, but that it is also significantly more robust when it comes to construing new possible counterexamples. Moreover, I will argue that the only promising way to attack P3_{g5} involves making an important concession to defenders of the argument from analogy – namely, the concession that particular experiences can simultaneously play both an enabling and an evidential role in a given instance of knowledge.

In order to reach the conclusion just announced, consider again case 3, which has served to refute P3_{g4}. As mentioned above, the case does not constitute a counterexample to P3_{g5} because there is some kind of exclusively evidential experience for Anna – the experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson – which does not carry the same evidential burden for Susan. This imbalance can be avoided in several ways.

First, the case might be reconstrued such that the testimonial experience in question is the only exclusively evidential experience figuring in Susan's knowledge, and thereby carries the same evidential burden as for Anna. However, the resulting case is one in which the experience of looking at the photos attached to the newspaper report does not play an evidential role anymore. (Even if the case were reconstrued such that this kind of evidence would figure in Susan's coming to understand the term crimson, this kind of experience would – not only from the perspective of opponents of the argument from analogy, but also from the perspective of defenders of the argument – not play a double-role, for the kind of reasons already mentioned above.) But then, the resulting case would no longer work as a counterexample to any version of the principle considered so far, for the unwelcome analogy-claim that the experience of looking at the photos attached to the newspaper report plays an evidential role for Anna could only be yielded in the light of a case which involves the kind of experience in question playing a clearly evidential role.

Second, the imbalance in question might be avoided by adding to Anna's evidential experience the same kind of newspaper-related experience that also plays an evidential role for Susan. The resulting case would then be one in which Anna does not only acquire the relevant piece of knowledge by relying on someone else's claim that the colour of the home jerseys of the 1986 football team of Washington State University is crimson, but also partly by reading the newspaper article and looking at the photos attached to it. That is, the kind of experience of looking at the photos attached to the article would then be had twice by Anna – first, as playing an enabling role in the course of coming to understand the term 'crimson', and later, as playing an exclusively evidential role in the course of acquiring the relevant piece of knowledge. (Note that the fact that some kind of experience plays an exclusively evidential role in some instance of knowledge is compatible with the same kind of experience also playing an enabling role in the same instance of knowledge, given that the kind of experience in question occurs more than once.) However, regardless of whether this modification avoids the imbalance in question, the resulting case does not constitute a counterexample to P3_{g5}. For even if no imbalance of the kind in question would keep P3_{g5} from yielding the analogy-claim that the experience of looking at the photos attached to the newspaper report plays an evidential role for Anna, this should no longer be regarded as an unwelcome result, because given the modification just described, the claim that the experience of looking at the photos attached to the newspaper report plays an evidential role for Anna is uncontroversially true.

Third, the problem just described could be avoided by – instead of adding to Anna's evidential experience the same kind of newspaper-related experience also had by Susan – adding to Anna's evidential experience any other kind of exclusively evidential experience. Given this modification, however, the fact that Anna now has a kind of exclusively evidential experience that does not also figure as exclusively evidential experience for Susan disqualifies the resulting version of the case from constituting a counterexample to P3_{g5}. Moreover, avoiding this fact by adding the same kind of evidential experience for Susan is of no help; for given that the experience of looking at the photos attached to the newspaper report carries some evidential burden for Susan, but not for Anna (– otherwise the resulting version of the case could not constitute a counterexample to any version of the principle –), it is impossible that both the testimonial experience and the new kind of experience in question carry the same evidential burden for Susan and Anna.

I take it that the above considerations already demonstrate that P3_{g5} is significantly more counterexample-proof than the four previous versions of the principle. However, I do not claim that P3_{g5} is true. Not only would such a claim be inadmissibly hasty in the light of the fact that several further possible variations of the case of Susan and Anna have not been tested above; certain observations in the following section will even make it appear quite plausible to regard P3_{g5} as false. What I *do* claim, however, is that at the present point of the discussion of possible true principles behind P3, defenders of the argument from analogy have the following reason to be optimistic. Even though not all possible modifications of the case have been considered, the above discussion has shown that any modification consisting in additions or subtractions of exclusively evidential experience fails to provide a counterexample to P3_{g5}. The remaining possible candidates for avoiding the imbalance in question, resulting from the evidential-burden-requirement formulated in P3_{g5}, are the following two. The first consists in adding to Anna's evidential experience some kind of experience which does not play an *exclusively* evidential role. This kind of additional evidential experience could be a candidate for preventing an imbalance regarding the evidential burden of the testimonial evidence in question, while not creating a new imbalance (since the evidential-burden-requirement only holds regarding exclusively evidential experience). The second remaining possible candidate for avoiding the imbalance in question consists in eliminating all exclusively evidential experience for Anna (thereby vacuously satisfying the evidential-burden-requirement), by modifying the case such that all of Anna's evidential experience also plays an enabling role. (I take it that hardly anyone would wish to claim that the case might be modified such that none of the experience figuring in it plays an evidential role.)

The crucial observation here is that both remaining options just described consist in modifying the case such that some experience plays a double-role in it – the same kind of double-role that, if the argument from analogy is sound, is played by some of the experience figuring in Norman's knowledge that all crimson things are red in the original version of this case. If it turned out that it is in principle possible for experience to play the relevant kind of double-role, then a first step would be made towards making plausible the conclusion of the argument from analogy.³³ Therefore, defenders of the

³³ This is not exactly true. A possible way for an experience-token to play a double-role that has not been mentioned so far might be by playing an enabling role in someone's knowledge of a conjunct of a complex truth and playing an evidential role in the same person's knowledge of a different conjunct of the

apriority of R1 might wish to deny that experience can in principle play a double-role. However, as the above considerations have shown, the only possibility in sight for construing counterexamples to $P3_{g5}$ is by accepting that experience can play the kind of double-role in question. Thus, since P3 is an instance of $P3_{g5}$ – so that if $P3_{g5}$ is true, P3 is true as well – it seems that opponents of P3 have to accept that experience can play the kind of double-role in question. (I assume here that a principle is true if there are no counterexamples to it.) This observation diminishes the argumentative space available to opponents of the argument from analogy.

Opponents of P3 may still object that the above considerations merely concern a specific kind of example, and that these considerations do therefore not provide a good reason to think that $P3_{g5}$ cannot be refuted by some different kind of example not involving a commitment to the claim that it is possible for experience to play a double-role. However, the prospects for this objection are rather dim, for the following reasons. First, the structure of any possible counterexample would have to be the same as in the cases already considered; that is, a successful counterexample would have to involve two instances of knowledge of the same truth, with a certain kind of experience figuring in both instances of knowledge and playing an evidential role in one of them, but not in the other. Examples with a different kind of structure could not serve to refute $P3_{g5}$, because applying $P3_{g5}$ to any other kind of example (for instance to one not involving two instances of knowledge of the same truth) would only yield vacuously true claims. Second, the most promising candidates for examples sufficiently differing – within the limits just described – from the kind of example considered above in order for the above considerations not to apply to them are those in which the relevant pieces of knowledge are not empirical truths, but rather arithmetical or logical ones. In order for such cases to qualify as counterexamples to $P3_{g5}$, each of them will have to involve two instances of knowledge of the respective truth, as well as some kind of experience figuring in both instances of knowledge and playing an evidential role in only one of them. Not only is it unclear whether examples of this kind can be convincingly construed when the relevant pieces of knowledge are arithmetical or logical truths, but – more importantly – even if

same complex truth, thereby playing both an enabling and an evidential role in the same instance of knowledge of a complex truth. Showing that experience can play this particular kind of double-role would, in contrast to what I have claimed above, *not* constitute a first step towards making plausible the conclusion of the argument from analogy. However, my considerations in favour of the claim that opponents of P3 are forced to accept that experience can play a double-role are not at all concerned with the particular way of playing a double-role just described. This and similar ways of playing a double-role may therefore simply be ignored here.

the latter can be achieved, and even if it can be achieved in a way such that $P3_{g5}$ would be refuted, defenders of $P3_{g5}$ may still restrict the principle such as to exclude arithmetical or logical truths. Of course, a restriction of $P3_{g5}$ should not go as far as to exclude R1, for the resulting version of the principle would then no longer be instantiated by P3. However, I take it that the burden of proof now lies with those who claim that not only $P3_{g5}$, but also all versions of it involving restrictions not excluding R1, can be refuted by examples without a commitment to the claim that experience can in principle play a double-role.

In the following section, I will provide some further, independent motivation for the claim that it is possible for experience to play a double-role, together with some further considerations in favour of P3.

4.2. The criterion for playing a double-role

In the previous section, I have argued that the only available way to construe a counterexample to $P3_{g5}$ would involve conceding that it is at least in principle possible for a given experience-token to play both an enabling and an evidential role in a given instance of knowledge. In the present section, I will describe yet another version of the kind of case considered above – one that can plausibly be regarded as one in which some experience plays a double-role, independently of the considerations presented in the preceding section. Furthermore, I will argue that there is no promising candidate for a general criterion in virtue of which the relevant experience in the case in question plays a double-role which does not also apply to the original version of the case of Norman.

As already stated in the previous section, one of the remaining options for providing a counterexample to $P3_{g5}$ consists in construing a case in which all of Anna's evidential experience also plays an enabling role for her. Since the only evidential experience for Anna is the testimonial experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson, modifying the case in the desired way just described would consist in integrating the testimonial experience in question into the experience that provides Anna with an understanding of the term 'crimson'. And indeed, there is a promising way of doing this, yielding a version of the case (case 4 in appendix B) which can – even

independently of any considerations regarding $P3_{g5}$ – plausibly be regarded as one in which experience plays a double-role.

In cases 1-3, Anna comes to understand the term by looking at the photos attached to the newspaper report and being told that the colour of the jerseys is one of the paradigm samples of crimson. Since she does not read the report or the caption, she does, at this occasion, not come to know that the colour of the home jerseys of the 1986 football team of Washington State University is crimson. Imagine, instead, that she comes to understand the term by looking at the photos, reading the caption stating that the photos depict the home jerseys of the 1986 football team of Washington State University, and being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson. Since by merely looking at the photos and reading the caption, which does not contain the term ‘crimson’, Anna would not have come to understand the term, the testimonial experience in question clearly plays an enabling role in this version of the case. By relying on the testimony in question, she does not only come to understand the term ‘crimson’, but she also comes to know that the colour of the home jerseys of the 1986 football team of Washington State University is crimson. Since it would be very hard to deny that some experience plays an evidential role in the given instance of knowledge, and since the testimonial experience in question is the most obvious candidate for playing this role, it is quite plausible to hold that the testimonial experience in question does not only play an enabling role, but also an evidential one.

I do not wish to settle the question whether case 4 constitutes a counterexample to $P3_{g5}$. For this to work, it would have to be the case that Anna’s experience of looking at the photos does not play an evidential role. I take it that, given the peculiar characteristics of the case, there is no clear initial answer to this question. Thus, the question of the truth or falsity of $P3_{g5}$ will remain open. Instead, I merely wish to hold that *some* of Anna’s experience in case 4 – namely, the testimonial experience in question – plays a double-role. In the remainder of the present section, I will argue that it is hard to accept this latter fact without also accepting that some experience plays a double-role in the original version of the case of Norman.

In order to provide a convincing case for the claim that only Anna’s testimonial experience in case 4, but not some experience in the original version of the case of Norman (henceforth simply ‘the case of Norman’), plays a double-role, one should be

able to formulate a criterion marking this difference. I will now consider possible candidates for such a criterion.

First, one might regard the distinguishing criterion in question to consist in the fact that Anna's testimonial experience in case 4 does not by itself provide her with an understanding of the term 'crimson'. Rather, it does so in connection with the experience of looking at the photos and of reading the caption. However, not only does this candidate for the distinguishing criterion in question appear quite ad hoc – since it bears no obvious connection to the question of evidentiality –, but it is also fulfilled by some of the experience in the case of Norman. For given any plausible individuation of kinds of experience according to which there are different kinds of experience providing for Anna's understanding of the term 'crimson' in case 4, there are also different kinds of experience providing for an understanding of the term 'crimson' in the case of Norman. That is, Norman does not come to understand the term 'crimson' merely by looking at crimson things, but by a connection of this kind of experience with some testimonial experience conveying the necessary information that the things he is looking at are of a colour called 'crimson'. Both kinds of experience, the testimonial one just described and the visual one, fulfil the criterion in question.

Second, one might suppose that the distinguishing criterion in question consists in the fact that Anna's testimonial experience in case 4 figures in an instance of knowledge of a contingent, rather than a necessary, truth. Obviously, this criterion is not fulfilled by any of the experience figuring in the case of Norman (or at least not by any experience *with respect* to Norman's knowledge of R1 – some of the experience figuring in Norman's knowledge of R1 might of course also figure in some of Norman's knowledge of contingent facts). However, this candidate for a distinguishing criterion is problematic for several reasons. First, it is quite uncontroversial that knowledge of all kinds of necessary truths can be acquired on the basis of evidential experience (say, by reading off the result of a calculation from the display of a calculator, in the case of arithmetical truths). But given that this fact is uncontroversial merely with respect to somewhat indirect, unparadigmatic ways of coming to know necessary truths, and given that Norman acquires his knowledge of R1 in a direct, paradigmatic way, this point may simply be neglected. Second, and more importantly, the candidate in question is clearly in need of supplementary reasons for the claim that paradigmatic ways of knowing necessary truths cannot be based on evidential experience. Attempts at providing such supplementary reasons might, most promisingly, be undertaken within the framework of

two-dimensional semantics. The idea would be to hold that sentences involving terms with both a primary and a secondary intension – for an explanation of this terminology, see section 2.3 above – express (at least) two distinct propositions, a primary and a secondary one (see Chalmers 1996, pp. 63f.). Making this distinction would allow one to hold that all necessary propositions are a priori, by holding that, while a sentence like ‘water is H₂O’ expresses a necessary as well as an a posteriori proposition, there is no single proposition that is both necessary and a posteriori. Rather, while the a posteriori proposition expressed is the one containing the primary intension of the term ‘water’, the necessary proposition expressed is the one containing the *secondary* intension of the term. And while the first of these two propositions, the primary one, is contingent and a posteriori, the secondary proposition is necessary and a priori. Alternatively, one might hold that instead of two distinct propositions, a single complex one is expressed, which is contingent a posteriori on one dimension and necessary a priori on another. (In his 2002, pp. 165f., Chalmers favours this version of the framework.)

However, in the light of my considerations in the previous chapters, the view just described is problematic for the following reasons. If primary and secondary intensions are conceived of as being sensitive to the difference between the ways in which Norman and most others come to understand the term ‘crimson’, then the resulting framework will be in conflict with P1, as already described in section 2.4. That is, the resulting framework is based on the claim that Norman’s knowledge about the microphysical surface structure of crimson things does not provide him with an understanding of the term ‘crimson’, a claim that, as already mentioned in section 2.4, comes without any supporting arguments and should therefore be regarded as begging the question against P1 and my considerations in favour of it presented in the previous chapter.

If, on the other hand, primary and secondary intensions are conceived of as not being sensitive to the difference between the ways in which Norman (in the modified version of the case) and most others come to understand the term ‘crimson’ (– such an insensitivity might occur if primary and secondary intensions are conceived of as functions from possible worlds to extensions, and if the usual way of understanding colour terms involves sufficient information about standard perceivers and circumstances –), then it is – especially in the light of my considerations in section 3.5 – not at all clear why Norman, in the modified version of the scenario, should not count as grasping the proposition expressed by ‘All crimson things are red’ (or, alternatively, the dimension of this proposition) which contains the primary intensions of the terms

‘crimson’ and ‘red’. And since he is, in this scenario, nevertheless not yet in a position to know this proposition, the latter, which is – given the usual candidates for primary intensions of colour terms, that is, those involving phenomenal properties – clearly necessary, may well count as a posteriori. Claiming otherwise would be in danger of begging the question against proponents of the argument from analogy.

Given these considerations, other supplementary reasons for the claim that paradigmatic ways of knowing necessary truths cannot be based on evidential experience should be offered. Thus, the burden of proof regarding whether the criterion in question is a promising candidate for marking the relevant kind of difference lies with opponents of P3.

Third, one might regard the distinguishing criterion in question to consist in the fact that Anna’s testimonial experience in case 4 figures as episodic memory in her knowledge, whereas Norman does not come to know R1 by relying on any particular memory of seeing something crimson, but merely on his *generic* memory of the superficial looks of crimson things. As already mentioned in section 1.2, the distinction between episodic and generic memory plays a central role in Williamson’s refusal to characterize any of the experience in Norman’s knowledge of R1, as well as any of the experience figuring in his knowledge of R2 (i.e., of the fact that all recent volumes of *Who’s Who* are red) as strictly evidential, and thereby in his refusal to characterize these instances of knowledge as a posteriori (if aposteriority is to be considered as a useful epistemological category). As I have also stated in section 1.2, given a definition of the term ‘strictly evidential’ according to which strict evidentiality can only be a feature of experience yielding episodic memory, it is not plausible not to characterize an instance of knowledge as a posteriori just because none of the experience figuring in it plays a strictly evidential role. And as mentioned in section 1.3, Williamson provides no reason to lend such weight to the distinction between experience yielding episodic memory and experience yielding purely generic memory. In the case of Norman’s knowledge of R2, it is *prima facie* plausible to characterize Norman’s experience of perceiving the colour of recent volumes of *Who’s Who* as evidential, regardless of whether it yields episodic or purely generic memories. Moreover, it is not at all infelicitous or inappropriate for Norman to answer the question ‘What is your *evidence* for believing that all recent volumes of *Who’s Who* are red?’ by saying ‘I saw them’, even though his memory of perceiving the volumes’ colour is purely generic. Thus, the burden of proof regarding

whether the distinction between episodic and generic memory is of significant relevance concerning the issues at hand lies with opponents of P3.

The above considerations do, of course, not constitute a definite case for the claim that there is no distinguishing criterion in virtue of which Anna's testimonial experience in case 4, but no experience in the case of Norman, plays a double-role. However, I take the above considerations to show that if there is such a criterion, it will not be easy to find. Thus, I take it that the burden of proof regarding whether some experience in the case of Norman plays a double-role lies with those who answer this question in the negative. I will now go on to propose a candidate for a criterion in virtue of which some of Anna's enabling experience in case 4 plays a double-role – a criterion which is also fulfilled regarding the case of Norman.

Obviously, Anna could have come to understand the term 'crimson' in ways not involving the experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson. Modifying the case by replacing the experiences providing for Anna's understanding of the term 'crimson' by other experiences also providing for this understanding may yield a case in which she is not in a position to acquire the relevant piece of knowledge. Something analogous holds regarding the original case of Norman. That is, modifying the case by replacing the experiences providing for Norman's understanding of the term 'crimson' by other experiences also providing for this understanding (e.g., by experiences providing him with knowledge about the microphysical surface properties of crimson things) may yield a case in which he is not in a position to acquire the relevant piece of knowledge. Such a case is exemplified by the modified version of it involving Mary. In contrast, take, for instance, cases 1-3, in which, as assumed in the previous section, Anna's enabling experience of seeing the photos attached to the newspaper report does not play an evidential role. Replacing, in these cases, the experience providing for Anna's understanding of the term 'crimson' by other experiences also providing for this understanding could *not* (without any further modifications) yield a case in which she is not in a position to acquire the relevant piece of knowledge, because after any such modification, the case will still be one in which she has the experience of being told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson (since this experience is exclusively evidential in the original version and thereby not affected by the kind of modification in question), and this experience puts

Anna in a position to acquire the relevant piece of knowledge, regardless of the way in which she has come to understand the term ‘crimson’.

As already announced above, I propose the feature just described as a candidate for a criterion in virtue of which some of the enabling experience in a given instance of knowledge plays a double-role. This proposal amounts to the following general claim:

(CDR) For any instance of knowledge, if modifying the case by replacing the experiences providing for the subject’s understanding of one of the relevant terms by other experiences also providing for this understanding may yield a case in which the subject is not in a position to acquire the relevant piece of knowledge, some of the enabling experience figuring in the respective case plays a double-role.

Why should CDR be regarded as true? Here is a possible rationale. If the antecedent of CDR is true, mere understanding of the relevant terms (together with some corresponding offline reasoning skills) is not what puts the subject in a position to acquire the relevant piece of knowledge. Rather, what puts the subject in a position to acquire the relevant piece of knowledge is her understanding the relevant terms *in a certain way*. And one might hold that this difference affects the role played by the respective experience. Given that experience plays an enabling role only insofar as it provides for the subject’s understanding of the relevant terms (and maybe some corresponding offline reasoning skills), and given that in the kind of case in question, experience provides for more than just for some replaceable way of understanding (plus corresponding offline reasoning skills), the role of some of the experience is more than just an enabling one.

Against this line of reasoning, it may be objected that in those cases fitting the description in the antecedent of CDR, fulfilling the enabling role is simply more demanding than in cases in which any way of understanding the relevant terms puts one in a position to acquire the relevant piece of knowledge. Thus, it may be held, the experience in question still plays an exclusively enabling role. However, it seems that by endorsing this kind of objection, one enters a slippery slope. For claiming that the enabling role is more demanding because more than just some replaceable way of understanding (plus corresponding offline skills) is demanded for putting the subject in a position to acquire the relevant piece of knowledge may seem to commit one to the

unwelcome claim that paradigm cases of a posteriori truth may be come to know with experience also merely playing a – relatively demanding – exclusively enabling role. It seems, for instance, that one would be committed to accepting that someone who learned the term ‘H₂O’ by being told that H₂O is the chemical structure of water, and who thereby comes to know that water is H₂O has acquired this knowledge with experience merely playing a – relatively demanding – exclusively enabling role, even though the subject’s being in a position to acquire the piece of knowledge in question does not merely depend on the subject’s understanding of the relevant terms, but on a specific way of understanding them. In order to avoid this commitment, one would have to provide an explanation of why in some, but not in all, cases described in the antecedent of CDR, the consequent of CDR holds (– an explanation of why the consequent of CDR holds regarding the case just described, but not regarding the case of Norman). As long as such an explanation has not been provided, the above-presented rationale constitutes a promising argument for CDR. And since the antecedent of CDR is fulfilled in the case of Norman (at least if P1 is true), the rationale for accepting CDR also constitutes a promising argument for P3.

In the present section, I have tried to shift the burden of proof to those who claim that there is a distinguishing feature in virtue of which some of Anna’s experience in case 4, but no experience in the case of Norman, plays a double-role; and I have proposed a candidate for a feature in virtue of which experience plays a double-role in both cases. However, it should be conceded that these considerations are far from constituting a definite case for P3. In the following section, I will therefore strengthen my case for P3 by presenting an additional argument in favour of it.

4.3. Another argument for P3

In the present section, I will provide a supplementary argument in favour of P3, that is, an argument to the conclusion that if the experience of seeing crimson things with the unaided eye plays an evidential role in Norman’s knowledge that all crimson things are red in the modified version of the case (the version involving Mary), then this kind of experience also plays an evidential role in Norman’s knowledge that all crimson things are red in the original version of the case. I call the argument in question the omega argument, due to a certain aspect of its structure (i.e., the aspect of containing a kind of

detour, symbolized by the shape of the Greek symbol 'Ω'). The first premise is the following:

(ΩP1) If the experience of seeing crimson things with the unaided eye plays an evidential role in Norman's knowledge that all crimson things are red in the modified version of the case, then the same holds regarding the experience of seeing *red* things with the unaided eye in this version of the case.

Here is a rationale for ΩP1. In the modified version of the case (the version involving Mary), as in the original version (the one construed by Williamson), Norman comes to know that all crimson things are red via offline reasoning, employing his visual memory of the superficial looks of crimson things and red things. Apart from the fact that the red-seeing experiences play an enabling role, while the crimson-seeing experiences may play no such role (since Norman, as argued in my case for P1, already understands the term 'crimson' before having these experiences), there is, regarding the issue of evidentiality, no significant difference between the way in which the two kinds of experience yield the respective offline skills, or between the ways in which the two offline skills are employed. (I take it that the fact that Norman comes to know that all crimson things are red, rather than that not all red things are crimson, is irrelevant here.) And as argued in the preceding two sections, opponents of P3 have to concede that the fact that some experience plays an enabling role in a given instance of knowledge does not prevent this experience from also playing an evidential role in this instance of knowledge. Thus, the fact that the red-seeing experiences play an enabling role does not prevent the similarities just mentioned between the roles of the red-seeing experiences and the crimson-seeing experiences from making it very plausible to assume that either none or both of these kinds of experience play an evidential role in the modified version of the case.

The next step in the argument is the following premise:

(ΩP2) If the experience of seeing red things with the unaided eye plays an evidential role in Norman's knowledge that all crimson things are red in the modified version of the case, then the same holds regarding this kind of experience in the original version of the case.

Given that the two versions of the case only differ with respect to the way in which Norman comes to understand the term ‘crimson’, it would be quite hard to deny that the red-seeing experiences play the same role in the two versions. That is, if they play a double-role in the modified version, then the same holds regarding the original version.

The final premise is the following:

(Ω P3) If the experience of seeing red things with the unaided eye plays an evidential role in Norman’s knowledge that all crimson things are red in the original version of the case, then the same holds regarding the experience of seeing *crimson* things with the unaided eye in this version of the case.

I take it to be obvious that if Ω P1 is true, Ω P3 is true as well. As mentioned above, the only significant difference, regarding the issue of evidentiality, between the roles of the crimson-seeing experiences and the red-seeing experiences in the modified version of the case is that the latter kind of experience plays an enabling role, while the former may not. In the original version of the case, not even *this* difference occurs, so that it would be very implausible to hold that there is any significant difference between the roles of the two kinds of experience in the original version of the case.

The three premises listed above entail, by the transitivity of the material conditional, the following conclusion:

(Ω C) If the experience of seeing crimson things with the unaided eye plays an evidential role in Norman’s knowledge that all crimson things are red in the modified version of the case, then the same holds regarding this kind of experience in the original version of the case.

Despite its relative simplicity, I take the omega argument, whose conclusion obviously entails P3, to significantly strengthen the case for P3 already provided in the preceding two sections. The most vulnerable part of the omega argument is clearly Ω P1, since it is the only premise involving an analogy between experiences which may differ regarding the issue of playing an enabling role. However, in the light of my considerations in favour of the claim that playing an enabling role does not prevent an experience-token from playing an evidential role in the same instance of knowledge, it is not clear how the difference in question could figure in a convincing argument against Ω P1.

Before closing, one possible source of motivation for rejecting the omega argument remains to be addressed. The kind of motivation in question is one that would derive from endorsing a certain version of two-dimensionalist semantics, namely one according to which the primary intension of a term may vary from one person to another. As already stated in section 2.4, the plausibility of this kind of view would not threaten the first half of the argument from analogy (that is, the argument to C1). However, as announced in section 2.5, the view in question may well be regarded as being in conflict with P3. For applying a subject-sensitive notion of primary intension might yield the following evaluation of the two cases involving Norman. While the primary intension of the term ‘crimson’, as used by Norman in the modified version of the case, involves some kind of reference to microphysical surface properties, the primary intension of the term ‘crimson’, as used by Norman in the original version of the case, involves some kind of reference to a certain phenomenal appearance to certain kinds of perceivers under certain kinds of circumstances. (As stated in section 2.3, this distinction might only yield a difference if intensions are not conceived of as functions from possible worlds to extensions, but rather as cognitively individuated entities.) Since in the modified version of the case, the knowledge that all crimson things are red cannot be acquired merely on the basis of the primary intensions of the relevant terms (the primary intension of ‘red’ is the same in both versions of the case), the crimson-seeing experiences, which finally put Norman in the position to acquire the piece of knowledge in question, plays an evidential role, and Norman’s knowledge – so the subject-sensitivist’s evaluation continues – is thereby a posteriori. In contrast, in the original version, the crimson-seeing experiences provide for Norman’s grasp of the primary intension of ‘crimson’, thereby playing an enabling role; and since the piece of knowledge in question can be acquired solely on the basis of the relevant primary intensions, Norman’s knowledge in this version of the case – so the opponent of P3 could conclude – is a priori. The following passage by Jackson & Chalmers, which has already been cited in chapter 2, indeed suggests the kind of treatment just described:

“[P]erhaps a city-dweller might use ‘water’ nondeferentially for the liquid that comes out of faucets (knowing nothing of oceans), and a beach-dweller might use ‘water’ nondeferentially for the liquid in the oceans (knowing nothing of faucets). This sort of case can be treated as we suggested that cases involving names be treated: the subjects have different conditional abilities, and different associated conditionals will be knowable a priori for each of them.” (Jackson & Chalmers 2001, p. 328)

If the kind of reasoning considered here were sound, then ΩC would be false. In particular, the kind of reasoning considered here is in conflict with $\Omega P1$. For proponents of this kind of reasoning would very likely claim that the red-seeing experiences play a merely enabling role in both versions of the case, since – in contrast to the crimson-seeing experiences in the modified version – they provide for Norman’s grasp of the primary intension of the respective term.

However, it is not at all clear why the kind of reasoning in question should be regarded as contributing anything substantially new to the discussion of the argument from analogy. For given the above-quoted description of the conception of primary intensions provided by Jackson & Chalmers, according to which using a term with a certain primary intension amounts to possessing certain conditional abilities, the claim that terms can be used with different primary intensions by different speakers seems to amount to something not significantly different from the claim that different speakers can come to understand terms in different ways. Accordingly, the fact that the red-seeing experiences in the modified version of the case provide for Norman’s grasp of the primary intension of the respective term would seem to amount to something not significantly different from the fact that he comes to understand the term ‘red’ in the usual way, that is, in a way providing him with the ability to detect red things by looking at them with the unaided eye. Thus, the facts that proponents of the above-described kind of reasoning can point to in attempting to refute $\Omega P1$ seem to amount to nothing substantially different from those facts that have already been considered in my arguments supporting the premises of the argument from analogy. In particular, the supposed fact that the red-seeing experiences provide for Norman’s grasp of (what counts, regarding him, as) the primary intension of the respective term would at most prove that they play an enabling role. But I have already argued that this would not prevent these experiences from also playing an evidential role in this case. Describing the cases in question in terms of primary intensions varying from one subject to another does therefore not seem to provide any new threat for defenders of the omega argument, or for defenders of the argument from analogy.

Of course, defenders of the kind of reasoning in question may try to provide a conception of primary intensions that avoids the point just made, that is, a conception that clearly distinguishes primary intensions from ways of understanding. However, in order to be of any use for attacking the omega argument (and also in order to avoid objections provided in sections 2.3 and 2.4), the new conception would still have to

allow for the possibility that primary intensions vary across speakers, as well as the possibility that a speaker understands two colour terms in different ways. And I take it to be, at least *prima facie*, unclear how a conception of this kind could avoid the point made in the preceding paragraph. Thus, the burden of proof lies with those who claim that the omega argument can be refuted on the basis of the kind of reasoning in question.

As already mentioned above, I take the omega argument to significantly strengthen the case for P3 already provided in the preceding two sections. And as just argued, it is not clear how endorsing a subject-sensitive version of two-dimensionalism could provide a good reason for revising this evaluation.

4.4. Summary and outlook

Having defended the first half of the argument from analogy in chapters 2 and 3, my aim in the present chapter has been to defend the remaining part of it. Since C1 and P3 obviously entail C2, the only part of the argument that remained to be defended was P3. In section 4.1, I have considered several candidates for a general principle instantiated by P3. In the course of these considerations, it has turned out that the only remaining possible kinds of counterexamples to a certain version of the kind of principle in question are such that some experience in them plays a double-role. Thus, it has turned out that there is either some true principle behind P3, in which case P3 will be true as well, or that it is in principle possible for some enabling experience to also play an evidential role in the same instance of knowledge. Since opponents of P3 will obviously be reluctant to accept the first of these two possibilities, and since the claim that experience can in principle play a double-role is a first step in arguing for P3, the considerations in section 4.1 diminish the argumentative space available to opponents of P3. In section 4.2, I have, by presenting a certain example, provided some independent motivation for accepting that experience can play a double-role. Furthermore, I have argued that there is no promising candidate in sight for a distinguishing criterion in virtue of which some experience plays a double-role in the example presented, but not also in the original version of the case of Norman. I have thereby shifted the burden of proof to opponents of P3. Moreover, I have provided some considerations in favour of a different candidate for a criterion in virtue of which some experience plays a double-

role in the example presented – a criterion according to which a double-role is also played by some experience in the original version of the case of Norman. Thus, the considerations in section 4.2 constitute a case for P3, though admittedly not a definitive one. Finally, in section 4.3, I have provided what I take to be a strong argument in favour of P3 – again, partly relying on the claim that experience can in principle play a double-role –, and I have defended this argument against a certain line of attack.

In sum, I take my considerations in this chapter to endow P3 with a considerable amount of plausibility. However, my defense of the argument from analogy is not yet complete. In the following chapter, I will consider a line of attack addressing a certain presupposition regarding colour terms that I have been making throughout the foregoing chapters.

5. Examining the argument from analogy, part IV: P1 in the light of arguments for colour relativism

In this chapter, I will question a presupposition regarding colours made throughout the foregoing chapters. In the course of that, I will examine a threat for defenders of the argument from analogy posed by certain arguments in favour of a relativist theory of the meanings of colour terms. In section 5.1, I will argue that a successful case for colour relativism would have to be regarded as undermining P1. In section 5.2, I will argue that the occurrence of so-called standard variation phenomena poses a serious threat for the most commonly endorsed versions of colour absolutism, and that a certain version of colour relativism appears to provide an elegant way of avoiding this threat. In sections 5.3 and 5.4, I will consider – and dismiss – various options for defending the argument from analogy against the threat posed by the occurrence of the phenomena in question. In section 5.5, I will argue that, contrary to how things may initially appear, the occurrence of standard variation phenomena poses a challenge for colour relativism not less serious than the one it poses for the most commonly endorsed versions of colour absolutism. Finally, in section 5.6, I will consider the remaining options for avoiding the challenge posed by the phenomena in question as well as the relation of these options to the argument from analogy.

5.1. Colour relativism as a threat to P1

Throughout the foregoing chapters, I have made an assumption which is crucial regarding the prospects of the argument from analogy, and which has been left unquestioned so far. According to my description of the modified version of Williamson's original scenario, presented in section 1.4 above, Norman is told by Mary all there is to know about the microscopic surface properties that all and only crimson things necessarily possess. The unquestioned assumption that I have in mind is presupposed in this part of the description of the scenario – it is the assumption that there *are* any microscopic surface properties necessarily possessed by all, and only by, crimson things. In the present chapter, I will finally question this assumption. Moreover, the considerations to be presented in the following sections do not only threaten the

assumption that there are any microscopic *surface* properties necessarily possessed by all, and only by, crimson things, but they equally threaten the possible fallback assumption that all and only crimson things necessarily emit, under certain normal circumstances, a particular kind of light. (It is, anyway, implausible to accept this fallback assumption without also accepting the analogous assumption concerning microscopic surface properties, due to supervenience considerations.) Also, as it should be expected, the considerations to be presented in the following sections do not only concern the colour crimson; rather, they are also supposed to undermine the assumption that for *some* colours, there are microscopic surface properties (or light emission properties) necessarily possessed by all, and only by, things of that colour. In what follows, I will abbreviate the negation of the latter assumption as follows:

(NHB) No colours have necessary and unique hidden bases.

NHB is obviously incompatible with my description of the scenario of Mary and Norman presented in section 1.4., because given NHB, it is (due to presupposition failure) not true that Mary tells Norman all there is to know about the microscopic surface properties that all and only crimson things necessarily possess. However, NHB does not rule out the possibility that there are metaphysically contingent covariations between colours and certain hidden properties of things of the respective colour – covariations which, even though they do not hold in every possible world, still hold in the actual world as well as in all close possible worlds. Assuming that Mary tells Norman all there is to know about these covariations, it is still plausible to hold that the information provided by Mary can provide Norman with the reliable skill of correctly applying the word ‘crimson’ on the basis of observing certain hidden properties. And since P1 is the claim that Norman understands the word ‘crimson’ in virtue of the possession of this reliable skill, a convincing case for NHB would not automatically constitute a threat to the argument from analogy. However, the considerations to be presented in the following sections are not only supposed to favour NHB; they are also supposed to favour the following claim:

(Rampant Variation) In the actual world, there is not even a rough distinguishing correlation between any colour and any hidden surface (or light emission) properties of things.

In contrast to NHB, the truth of Rampant Variation would pose an undeniable threat to the argument from analogy, since no matter what Mary tells Norman about microscopic surface properties (or light emission properties) of objects, this kind of information will – given Rampant Variation – not put Norman in a position to reliably distinguish crimson things from things of different colours. So unless the scenario of Mary and Norman is substantially modified, so as to provide a different explanation of how Norman acquires the reliable skill of distinguishing crimson from non-crimson things, the truth of Rampant Variation would undermine P1.

There is a close connection between Rampant Variation and a certain view about colour that goes under the name ‘colour relativism’. This latter view consists in the claim that which colour an object has at a certain time is not an absolute matter. Rather, objects have certain colours at certain times only relative to certain perceivers under certain circumstances. This view has mainly been defended in Jackson 1996 and 1998, McLaughlin 2003a and 2003b, and Cohen 2004 and 2009. In order to get an initial grasp of the implications of this view, consider how the phenomenal appearance of an object that phenomenally appears green to most humans under normal daylight conditions – and which would therefore, in non-relativist terms, simply be regarded as being green – may vary together with the lighting conditions. Under dim moonlight conditions, for example, most humans are unable to distinguish green things from red things (speaking, for the sake of brevity, in absolutist terms) on the basis of their phenomenal colour appearance. Under this kind of lighting condition, both green and red things phenomenally appear dark grey to most humans. This kind of scenario does of course not prevent ordinary speakers who have observed the objects in question both under daylight and under dim moonlight conditions from claiming that the objects are green and red, respectively – not just green and red relative to certain conditions, but simply green and red. That is, in ordinary speakers’ attributions of colour, certain viewing circumstances are given a privileged role. Unsurprisingly, the viewing circumstances which play this privileged role in ordinary speakers’ attributions of colour are those to which our perceptual apparatus has been evolutionarily adjusted, and which are therefore more interesting to us than viewing circumstances under which our ability to distinguish colours is significantly impaired. Accordingly, Jackson – in making a case for relativism – has to make the following concession:

“[S]eeing something from somewhere between a third of a metre and ten metres, in daylight during most parts of the day, or in typical indoor lighting, are all good for detecting the shape, distance away, size and relative position of the objects around us, and it is the colour of objects in such normal circumstances that especially interests us.” (Jackson 1998, p. 97)

Thus, we as humans have good reasons to regard certain viewing conditions as normal or standard ones. And these reasons may also provide a motivation for regarding the viewing conditions in question as playing a privileged role in determining the colours of objects. Accordingly, these considerations may provide a motivation for attributing the colours green and red, respectively, in an absolute sense, to the two objects mentioned above, since this characterization supposedly corresponds to the objects’ respective phenomenal appearance to normal human beings under normal circumstances.

However, in the relativist framework, this way of characterizing the two objects is “thoroughly anthropocentric” (Jackson 1998, p. 97), and “the fundamental notion is that of the colour of [object] *O* at *T* for *S* in *C*” (ibid.), where ‘S’ and ‘C’ may also pick out subjects and circumstances that are not regarded as normal or standard at all from an ordinary human speaker’s perspective. Thus, in a relativist framework, the two objects mentioned above do not count as green and red in an absolute sense, but merely as green and red for certain perceivers under certain lighting conditions. Moreover, in this kind of framework, the two objects equally have the property of being dark grey for certain perceivers under certain circumstances – e.g., for normal human beings in dim moonlight –, and none of these ways of attributing colours to the two objects in question is more fundamental than the other.

The aforementioned close connection between colour relativism and Rampant Variation is now apparent. For present purposes, it is merely necessary to point out that colour relativism, together with the indisputable fact that the kind of variation of phenomenal appearance described in the example of the two objects given above is an actual common phenomenon, leaves no room for denying Rampant Variation. For given that any object will, in a relativist framework, count as being of different colours, relative to sufficiently diverging kinds of perceivers and circumstances, and that any two objects will count as being of the *same* colour, relative to sufficiently diverging kinds of perceivers and circumstances, it is clear that there can be no reliably distinguishing correlation between any colour and any intrinsic physical property.

It should be noted, however, that this diagnosis concerning the connection between colour relativism and Rampant Variation is, in a sense, over-simplified. Consider the

following more specific version of Rampant Variation, which is entailed by the more general one:

(Rampant Variation_{crimson}) In the actual world, there is not even a rough distinguishing correlation between crimson and any hidden surface (or light emission) properties of things.

I have characterized colour relativism as the view that there are only relativized, and no non-relativized, colour properties, and these properties can be referred to by uses of colour terms involving explicit relativizations to kinds of perceivers and circumstances. From this, it follows that if uses of colour terms *not* involving explicit relativizations to kinds of perceivers and circumstances are to refer to any property at all, they should contain *implicit* relativizations of the relevant kind. And as will be seen below in the description of the case for colour relativism, the phenomena motivating the claim that there are these implicit relativizations also motivate, to the same extent, the claim that the relata in question can vary across contexts of utterance; this context-variance claim might therefore simply be regarded as being part of the view, even though it does not explicitly appear in the formulations of the view provided by its main proponents mentioned above.

Thus, if colour relativism is correct, the semantic value of the term ‘crimson’ will, for each use of the above formulation of Rampant Variation_{crimson}, involve a certain kind of relativization to perceivers and circumstances. And if this relativization is specific enough (involving, say, perceivers relevantly similar to the respective speaker as well as normal daylight conditions), the proposition expressed (or whatever is regarded as the primary truth-value-bearer) by the respective use of the above formulation of Rampant Variation_{crimson} might well be false. And since the falsity of Rampant Variation_{crimson} would, at least given fixed implicit relata, entail the falsity of Rampant Variation, my above claim concerning the connection between colour relativism and Rampant Variation – that is, the claim that colour relativism, together with certain indisputable facts regarding the actual occurrence of phenomenal colour variation, leaves no room for denying Rampant Variation – might well be false for a certain range of possible implicit relata. In order to fix this, I stipulate that Rampant Variation is to be interpreted as being relativized to the widest possible range of perceivers and circumstances, that is,

to all possible perceivers and circumstances (whatever kind of possibility this will most appropriately involve).

Given this unrestricted interpretation of Rampant Variation, my above claim about the connection between colour relativism and Rampant Variation is vindicated. For given colour relativism and the kind of indisputable facts about phenomenal colour variation mentioned above, *no* object in the actual world would be of any colour relative to all possible perceivers and circumstances. Accordingly, there would be no *distinguishing* correlation between any colour and any hidden surface (or light emission) properties of things, because the correlations between any colour and any hidden surface (or light emission) properties of things would be exactly the same (since they would all be vacuous).

In my above explanation of why the truth of Rampant Variation would undermine P1, I pointed out that no matter what Mary tells Norman about microscopic surface properties (or light emission properties) of objects, this kind of information will – given Rampant Variation – not put Norman in a position to reliably distinguish crimson things from things of different colours. Now, for this latter claim to be convincing, the term ‘crimson’ (as well as the term ‘colours’) is to be interpreted in the same unrestricted way as Rampant Variation, that is, as relativized to all possible perceivers and circumstances.³⁴ Accordingly, if P1 is to be undermined by the unrestricted version of Rampant Variation, it has to be relativized analogously. The corresponding explicit formulation would be the following:

(P1_{unres}) In virtue of possessing the reliable skill of distinguishing those objects that are crimson *relative to all possible perceivers and circumstances* from those objects that have a different colour *relative to all possible perceivers and circumstances* on the basis of observing the surfaces of

³⁴ Given colour relativism, people who – unlike Norman – have acquired the word ‘crimson’ in a normal way have also not thereby been put in a position to reliably distinguish those objects that are crimson relative to all possible perceivers and circumstances from those objects that have a different colour relative to all possible perceivers and circumstances – again, because there would, given colour relativism, be no *distinguishing* correlation between any particular colour and the colour-indicators used by these people. But for the claim that colour relativism would undermine P1 to be promising, its truth should better not have the consequence that, given colour relativism, *no* speaker understands any colour term. However, it may well be argued that there is no such consequence, since one might hold that given colour relativism, a certain way of associating colour terms with certain phenomenal qualities is already sufficient for understanding the respective terms.

the objects in question through a microscope, Norman understands the word ‘crimson’.

One further clarificatory remark regarding the connection between colour relativism and P1 should be made here. According to the way in which I have described the issue, the truth of colour relativism would undermine $P1_{unres}$ because, contrary to what is stated in $P1_{unres}$, observing surfaces of objects through a microscope would not *provide* Norman with a basis for reliably singling out those objects that are crimson relative to all possible perceivers and circumstances, since there would be no distinguishing correlation between the microphysical properties in question and colours relativized to all possible perceivers and circumstances. However, assuming that colour relativism is true and that Norman comes to know this, he will, given some unsurprising empirical information regarding possibilities of phenomenal colour variation, be in a position to know that no object is crimson relative to all possible perceivers and circumstances. He is thereby also in a position to reliably distinguish those things that are crimson relative to all possible perceivers and circumstances from things that are not – simply due to the vacuousness of the task. And – in accordance with the aim of the original scenario of Mary and Norman – he is in a position to reliably fulfil the task in question without associating any phenomenal colour appearance with the term ‘crimson’. However, attempting to reconstruct $P1_{unres}$ along such lines, in order to vindicate it in the face of colour relativism, would not be promising at all, since it would be utterly implausible to claim that Norman understands the word ‘crimson’ in virtue of possessing the reliable skill in question in the vacuous manner just described. In particular, he would not be in a position to conceive of any relevant difference between crimson and any other colour that he does not associate with certain phenomenal qualities. There thus seems to be no way to vindicate $P1_{unres}$ in the face of colour relativism.

Obviously, the above remarks about the connection between colour relativism and P1 already invite for a consideration of possible different versions of P1 (among those, of course, versions involving more specific relata) which would not be undermined by the truth of colour relativism. I will consider the prospects of several such versions later on, in section 5.4. But first, I will examine the prospects of $P1_{unres}$ by examining the prospects of colour relativism.

5.2. The case for colour relativism

In this section, I will look at pieces of evidence for the claim that significant variation of phenomenal colour experience does not only occur when so-called standard lighting conditions, e.g. daylight or good lamplight conditions, turn into bad, and therefore non-standard, conditions, such as dim moonlight. Rather, the phenomena to be considered are supposed to be evidence for the claim that significant variation of phenomenal colour appearance occurs within the range of viewing conditions commonly regarded as standard, and for perceivers whose colour perception, in the circumstances in question, would as well commonly be regarded as standard. The problem supposedly constituted by this kind of phenomena is generally called the problem of standard variation. Later on in this section, I will describe how embracing colour relativism might be regarded as the most plausible way to deal with this problem, and I will address a few objections that defenders of the view can convincingly refute.

It may be worth noting that the problem of standard variation is relatively independent of debates about the metaphysical nature of colours. More precisely, the problem equally arises for any theory regarding the metaphysical nature of colours according to which objects are coloured in a certain way in virtue of phenomenally appearing in a certain way to standard perceivers under certain viewing conditions – be it a theory according to which colours are dispositions to cause certain phenomenal colour impressions for standard perceivers under standard viewing conditions, or a theory according to which colours are the physical bases of these dispositions, or further possible theories according to which colours are second-order properties of having one of the two properties just mentioned. What makes appearance variation for standard perceivers under standard conditions problematic for a colour absolutist who holds a theory of the kind just described is that it forces her (at least if she wishes to avoid the eliminativist conclusion that the respective objects are not really coloured at all) to characterize some of the perceptual states and/or viewing conditions that are commonly regarded as standard as not relevant, that is, as not standard enough, for determining the real colour of the object in question. And depending on how rampant the problem of standard variation turns out to be, such a characterization may involve an unacceptable degree of arbitrariness. In order to evaluate this degree, I will now look at the supposedly problematic phenomena.

As already mentioned, colour relativism has mainly been defended in Jackson 1996 and 1998, McLaughlin 2003a and 2003b, and Cohen 2004 and 2009.³⁵ Jackson, however, gives a relatively sparse presentation of cases of standard variation. In fact, he merely offers the following example:

“[T]he colored patches in many magazines look red from normal viewing distances but are revealed as made up of small magenta and yellow dots on closer inspection.” (Jackson 1996, p. 205, & Jackson 1998, p. 95)

The particular relevance of this kind of example consists in the fact that, while virtually *no* coloured object will phenomenally appear to be of the same kind of colour if observed through a microscope, the phenomenal disappearance of redness in the case of the magenta-and-yellow-dotted patches in magazines occurs much sooner in the process of looking more closely. That is, the phenomenal appearance of the patches as being made up of magenta and yellow dots is one that occurs to the unaided human eye. Therefore, choosing one of the two viewing distances as the one that reveals the real colour of the patches might be regarded as arbitrary. Accordingly, Jackson claims that

“it is plausible that both the looking red from a normal viewing distance and the looking made-up-of-yellow-and-magenta-dots from close up are color experiences that count as presentations of features of what is seen” (Jackson 1996, p. 205, & Jackson 1998, p. 96).

Moreover, by moderately increasing the size of the dots one can create patches regarding which it is even more plausible to claim that both the phenomenal appearance of red and the phenomenal appearance of magenta and yellow occur within a range of viewing distances that are commonly regarded as revealing an object’s real colour.³⁶

³⁵ Characterizing Cohen 2004 and 2009 as defenses of colour relativism is, however, not a straightforward matter – not only because he uses the label ‘relationalism’ instead of ‘relativism’, but also because he explicitly regards these two positions (the latter of which he ascribes to McLaughlin) as “importantly different” (Cohen 2009, p. 10). However, due to the fact that Cohen merely works with what he himself calls a “pretheoretical grasp of the relational/nonrelational distinction” (2004, p. 462, and 2009, p. 9), the exact differences between his relationalist view and Jackson’s and McLaughlin’s relativist view remain unclear. Several of Cohen’s remarks suggest that he regards his relationalism as a specific version of relativism. However, his characterization of relationalism as the view that “colors are constituted in terms of relations to visual systems and viewing circumstances” (2004, p. 471) may raise doubts regarding any significant difference between the two views. Cohen’s position has several times been identified with McLaughlin’s and Jackson’s relativism (see, for example, Kalderon 2007, p. 570), and I think it is at least safe to say that his arguments for colour relationalism also amount to arguments for colour relativism (whether or not the former is identical to or just a specific version of the latter).

³⁶ Those who are sceptical regarding the supposed fact that patches of magenta and yellow dots look red from an appropriate viewing distance may simply think of a more easily imaginable version of the example, for instance one involving green-looking patches which are in fact made up of blue and yellow dots. However, the crucial point of the original version of the example is not that magenta-and-yellow-dotted patches phenomenally appear *red* from an appropriate distance, but simply that they appear to have

As described above, the colour relativist has no problem in accommodating this kind of appearance variation. According to her theory, the patches are indeed red – but merely red for certain perceivers under certain viewing conditions, where these latter conditions include the viewing *distance*. And this claim is perfectly compatible with the relativist's further claim that the patches are magenta-and-yellow-dotted, and thereby not red, for certain perceivers under certain other viewing conditions, including a different viewing distance.

To the absolutist, there are different possible strategies available for accommodating the example. She might, first, concede that both viewing distances fall within the range of distances relevant for determining an object's real colour, and she might claim that, since an appeal to this range of distances does, in the example at hand, not purport to determine one single way of being coloured, the patches have no colour at all, that is, they are neither red nor magenta-and-yellow. This locally eliminativist strategy, however, involves some rather unattractive commitments. Consider, for example, a patch half the size of which is made up of smaller magenta patches and the other half of which is made up of smaller yellow patches. Moreover, imagine that the smaller magenta patches and the smaller yellow patches are still large enough, and arranged in such a way, that the whole patch phenomenally appears (to standard perceivers) as made up of magenta and yellow patches from any viewing distance commonly regarded as normal or standard. The strategy just mentioned involves a commitment to accept the very implausible claim that through a mere rearrangement of the positions of the magenta and yellow portions within the whole patch – that is, through rearranging them such that the whole patch will be magenta-and-yellow-dotted as in Jackson's magazine example – the patch stops being magenta-and-yellow, even though the new arrangement of the portions of the two colours is still coarse-grained enough so as to be detectable by the unaided eye.

A second possible strategy for accommodating the example within an absolutist theory simply consists in claiming that, since both viewing distances fall within the range of distances relevant for determining an object's real colour, the kind of patches in question are both red all over and magenta-and-yellow all over. One might initially think that endorsing this strategy would commit one to denying the truth of certain claims that most defenders of the a priori consider as paradigm a priori truths, such as

one uniform colour from an appropriate distance, be it red or any other. And this latter fact should appear quite plausible even to those who have never actually encountered patches of the relevant kind.

the claim that no surface can be uniformly red and uniformly blue at the same time (listed, for example, in BonJour 2005, p. 100, as one of five supposedly clear examples of propositions knowable a priori). However, if the above considerations lead one to accept that a surface can be both red and magenta-and-yellow at the same time, then there is still room for denying the further claim that any object whose surface is of some colour (in an absolute sense) can at the same time also be of any other colour. For accepting that a surface can both phenomenally appear red for normal perceivers under certain normal viewing circumstances and phenomenally appear yellow-and-magenta for normal perceivers under certain other, still normal, viewing circumstances is well compatible with claiming that if something both phenomenally appears red for certain perceivers under certain viewing conditions and phenomenally appears *blue* for certain perceivers under certain other viewing conditions, then either the perceiver or the viewing circumstances will be such that they should not be regarded as normal or standard. (It is, after all, hard to imagine that an example analogous to the one involving the magazine patches can be found in which the relevant pair of colours is not red and magenta-and-yellow, but red and blue.) Therefore, the absolutist strategy considered here does, at least *prima facie*, not seem unpromising.

A third possible absolutist strategy for accommodating the example consists in claiming that either the viewing circumstances under which the patch phenomenally appears red to normal perceivers or the viewing circumstances under which the patch phenomenally appears magenta-and-yellow to normal perceivers should not be regarded as revealing the patch's real colour. One might, for example, claim that within a certain range of viewing distances, shorter viewing distances should be given preference over longer ones when it comes to determining which of the distances, in conflict cases like the one under consideration, reveals the object's real colour. In the case of the present example, endorsing this version of the strategy would lead one to claim that the patch is magenta-and-yellow, and not red – at least if one regards the short viewing distance from which the patch looks magenta-and-yellow to the unaided eye of normal perceivers as falling within the relevant range of distances. Like the second strategy considered above, this third possible absolutist strategy for accommodating the example does, at least *prima facie*, not seem unpromising. It is therefore far from obvious why Jackson's example should be regarded as posing a serious threat to absolutist theories of colour. Moreover, he seems to be willing to admit that standard variation phenomena merely amount to examples of this kind: “[M]ostly objects look much the same colour

in all the circumstances we count as normal. The aforementioned coloured patches are something of an exception.” (Jackson 1996, p. 206, & 1998, p. 97) It can therefore be concluded that the problem of standard variation, as presented by Jackson, does not seem particularly threatening.

Both McLaughlin and Cohen, however, mention a vast range of further examples of standard variation phenomena. These examples significantly differ in several respects from the one presented by Jackson. One of these respects concerns the degree of generality of the colour properties in question. That is, the examples mentioned by McLaughlin and Cohen mostly concern relatively specific *shades* of red, blue, etc., which may – as McLaughlin and Cohen implicitly claim – also, despite their high degree of specificity, be regarded as proper colours. Unsurprisingly, the more specific a colour property, the more sensitive its attributions to small changes of lighting conditions. Accordingly, McLaughlin claims that even within the range of lighting conditions ordinarily considered as normal, one single object can phenomenally appear to be of different specific colours:

“The dress may look one shade of unique blue in the lighting of the shop, another in the bright sunlight of the street, and yet another by the light inside the restaurant, even though each lighting condition falls within what counts as the normal range by actual ordinary standards.” (McLaughlin 2003a, pp. 117f.)

And for those who hold that daylight conditions should be given preference over lamplight conditions regarding the question which type of condition reveals an object’s real colour, the problem of standard variation concerning specific colour properties – again unsurprisingly – reappears:

“As lighting conditions change over the course of a typical sunny summer afternoon, a particular tomato may look, perhaps, a dozen different basic shades [that is, shades of which there are no more determinate shades] of red due to shifts in lighting that fall within the normal range.” (McLaughlin 2003a, p. 117)

Furthermore, as pointed out in an early discussion of the problem of standard variation by Bertrand Russell, colour appearance varies according to the angle from which an object is perceived:

“[I]f several people are looking at the table at the same moment, no two of them will see exactly the same distribution of colours, because no two can see it from exactly the same point of view, and any change in the point of view makes some change in the way the light is reflected.” (Russell 1912, pp. 11f.)

And even if a non-arbitrary decision could be made in favour of one kind of lighting condition and viewing angle specific enough to rule out the kind of effects just mentioned, the problem of standard variation still arises, due to the different possible states of adaptation of the eye:

“Coming in from the sunlight, the protopigments will be bleached, and so the chair might look pinkish until one's eyes adjust to the different light, at which point it looks more reddish, yet both states of adaptation fall within the normal range.” (McLaughlin 2003a, p. 118)

Another aspect of the problem of standard variation concerns collocation effects, that is, intra-subject variations of phenomenal colour appearance of an object due to changes of the visual background against which the object is observed (cf. McLaughlin 2003a, pp. 118f., and Cohen 2004, pp. 454ff.). And as Block 1999 describes (cf. pp. 41-43), referring to studies presented in Neitz & Neitz 1998, Neitz, Neitz & Jacobs 1993, and Neitz & Jacobs 1986, there are even reliably testable cross-subject variations in colour perception corresponding to differences in gender, skin pigmentation, and age, due to differences in cone sensitivities of the respective groups. As Block puts it, privileging the colour perception of people of a certain gender, skin pigmentation, and age, while regarding the colour perception of the other groups as non-standard “would be sexism, racism, or ageism” (Block 1999, p. 43).

Taken together all the different aspects of the problem of standard variation, it is indeed undeniable that, in the case of relatively specific colour properties, the same object phenomenally appears to be of different colours for normal perceivers and under normal viewing conditions, given any plausible conception of normality. Accordingly, McLaughlin provides the following diagnosis:

“The problem, then, is that for those ways things look in colour, there will be no physical property that disposes its bearers to look that way in colour to *all* perceivers that count as normal in *every* circumstance that counts as normal by actual everyday standards.” (McLaughlin 2003b, p. 491)

McLaughlin's diagnosis might suggest that the problem of standard variation merely concerns colour physicalism, or, more specifically, the view that colours of objects can be identified with physical bases of certain dispositions of those objects. However, as mentioned above, the problem equally concerns *any* theory regarding the metaphysical nature of colours according to which objects are coloured in a certain way in virtue of phenomenally appearing in a certain way to standard perceivers under standard viewing conditions, including not just theories according to which colours are physical bases of

certain dispositions, but also theories according to which, for instance, colours can be identified with the dispositions in question. This is due to the fact that, given the phenomena mentioned above, and considering specific colour shades as proper colours, it is not only the case that there is no physical property that disposes its bearers to phenomenally appear to have the same colour for all normal perceivers under all normal circumstances, but it is also the case that no object has a *disposition* (or a second-order property of having a disposition) to phenomenally appear to have a certain colour to all normal perceivers under all normal viewing circumstances. The reason why McLaughlin, and Jackson as well, merely present standard variation phenomena as a problem for the view that colours are physical bases of certain dispositions is – at least according to a charitable reading – not that they think that these phenomena pose no problem for colour dispositionalism, but rather that these authors have independent reasons for preferring the described version of colour physicalism over colour dispositionalism, and are therefore only interested in standard variation phenomena insofar as they pose a threat to the former kind of view.

The fact that there are no dispositions to, or physical properties that dispose their bearers to, phenomenally appear to have one determinate specific colour for all normal perceivers under all normal circumstances might, of course, also be taken to favour the view that there are no colours. Such a case for colour irrealism is, for instance, presented in Hardin 1988. According to McLaughlin, this kind of conclusion can be avoided:

“Rather than embracing colour irrealism, we should handle the problem by radically relativizing the colours; that is, we should relativize them to kinds of perceivers and circumstances so specific as to leave no room for variation in colour appearance.” (McLaughlin 2003a, p. 122)

If relatively specific colour shades, and even *maximally* specific ones (so-called basic shades), are regarded as proper colours, it seems that precluding variation in phenomenal colour appearance by relativizing colours to kinds of perceivers and circumstances would only work if the relata in question were also relatively – or, at least in the case of basic shades – maximally specific. And indeed, McLaughlin holds that the kind of relativization necessary for accommodating the perceptual variation phenomena mentioned above involves relata of a completely specific kind (Jackson and Cohen are less explicit regarding this issue):

“[I]t seems appropriate to say that how things look in colour to perceivers of a certain completely specific kind in circumstances of a certain completely specific kind is how things are in colour *for them in such circumstances*. This is a radical relativization of colour. It seems to me the right response to the problem of standard variation.” (McLaughlin 2003b, p. 493)

The kind of relativization McLaughlin has in mind will practically not differ much from conceiving colours as relative to a certain perceiver at a certain time *t*. However, it is not at all clear whether the perceptual variation phenomena mentioned above force the relativist to claim that *all* colours should be conceived as relative to completely specific kinds of perceivers and circumstances. It rather seems that the relata merely have to be as specific as necessary in order to rule out the possibility of standard variation regarding the colour in question; and in the case of relatively general colours, like red and blue, the variation phenomena mentioned above do *not* show that a relativization to completely specific perceivers and circumstances is necessary in order to rule out the possibility of standard variation. Given the empirical evidence mentioned so far, it therefore seems that the most appropriate relativist account would be one that is more modest than the one proposed by McLaughlin – namely one according to which the specificity of the relata proportionally varies with the specificity of the respective colour.

This version of the view is also favoured by considerations regarding the mechanisms of determining the relata. As Cohen plausibly holds, “the presuppositions of ordinary thought and talk about color tacitly provide us with values for these parameters“ (Cohen 2004, p. 471). That is, the specificity of the relata of colour properties expressed in a conversation or figuring in thought depend on presuppositions of the respective speakers or thinkers regarding the kind of perceivers and circumstances in relation to which the respective colour ascription has the same truth value. And in the case of ascriptions of relatively general colours, like red and blue, it is very plausible to suppose that even those thinkers and speakers who are informed about the variation phenomena mentioned above do *not* presuppose that their respective colour judgments or ascriptions only reliably hold relative to a completely specific kind of perceiver and a completely specific kind of circumstance. Rather, in line with the version of the view mentioned above, it is plausible to assume that the presuppositions vary in accordance with the specificity of the ascribed colour.

There is a connection between the preceding remarks and a possible worry according to which colour relativists are committed to denying the plausible claim that

speakers sometimes genuinely disagree about the colours of objects. This kind of worry might indeed be appropriate with respect to versions of colour relativism according to which the kind of perceiver to which colour ascriptions are relativized is always specific enough so as to actually exclude anyone except the speaker herself. However, the worry is not appropriate with respect to the modest version of the view mentioned above, since according to this version, the kind of perceiver to which colour ascriptions are relativized *does* – at least in the case of relatively general colours – often include other speakers. Accordingly, Cohen has no trouble in refuting the objection:

“Speakers can agree and disagree about the colors of objects because the visual systems and viewing circumstances to which ordinary color attributions are tacitly relativized are not individuated so finely that they are unrepeatably singletons, but include the visual systems of other human beings like ourselves and the viewing conditions they typically encounter.” (Cohen 2004, p. 473)

According to a related worry (brought forward, to some extent, by Byrne & Hilbert 2003, pp. 57f.), colour relativism does not, or at least not often enough, allow for the possibility of colour illusion – a possibility, so the objection goes, that any plausible view about colour should allow for to some extent. However, this worry might at most be appropriate with respect to versions of colour relativism according to which the kinds of perceivers and circumstances to which colour ascriptions are relativized always include the actual speaker and the actual circumstances the speaker finds herself in while making the respective colour judgment or ascription. But with respect to versions of colour relativism according to which the kinds of perceivers and circumstances to which colour ascriptions are relativized do not have to include the actual speaker and the actual circumstances, the objection fails for the following reason (see also Cohen 2004, pp. 473f.). A speaker might presuppose that the viewing conditions under which she actually finds herself while trying to identify the colour of an object in front of her are of a certain kind *vc1*, say, good daylight conditions. Her judgment that the object in question is, say, purple, is thus relativized to *vc1*, that is, she judges that the object in front of her is purple to certain kinds of perceivers under good daylight conditions. She might, however, be wrong about the kind of viewing conditions she finds herself in. That is, it might, for example, be the case that the object is, unbeknownst to her, partially illuminated by a red spotlight. So relative to the kind of conditions her judgment is relativized to, namely, good daylight conditions not involving any additional spotlight, (and relative to the kind of perceivers the judgment is relativized to) the object is not purple, but blue. The speaker’s colour judgment is therefore false,

and this error is due to the illusory impression that the object in question is merely illuminated by daylight. The speaker might also be under an illusory impression regarding the functioning of her own eyesight, which provides yet another possible source of colour illusion.

Apart from colour illusion, colour relativism of the kind favoured by Cohen also allows for different kinds of error regarding an object's colour. Seeing an object at night, under viewing conditions obviously insufficient for determining the object's colour, one might, relying on one's visual memory, judge that the object is green, making a judgment relativized to normal human beings and good daylight conditions. Even though the speaker is under no relevant illusion regarding the viewing conditions that she finds herself in while making the judgment, or regarding the functioning of her eyesight, the judgment might still be wrong – namely, if she simply misremembers how the object has phenomenally appeared to her under the kind of conditions her ascription is relativized to. This kind of error occurs even more often in cases of colour judgments in which the object in question is not actually present. One might therefore conclude that objections according to which colour relativism does not sufficiently allow for the possibility of error, or illusion, in particular, do not pose a serious threat to the view.

According to yet another related objection, brought forward by Byrne & Hilbert, a certain, particularly plausible, kind of explanation of a certain kind of behaviour is not available to colour relativists. The kind of explanation in question is exemplified by the following case:

“Imagine that you have just eaten a tasty crimson fruit, and that you are now looking at another fruit of the same kind. (To avoid irrelevant distractions about color language, imagine you are an Old World monkey.) You recognize the fruit as having the same distinctive shade of red as the first, and that's why you reach for it.” (Byrne & Hilbert 2003, p. 58)

As already indicated, Byrne & Hilbert claim that the kind of explanation of the fruit-eating behaviour given in the last sentence of the cited passage is one that colour relativists cannot account for. Their reason for this claim concerns the relativization to viewing conditions. They argue as follows:

“Call the first ‘type of circumstance of visual observation’ [...] C_{F1} , and call the second C_{F2} . Unless the relativization to types of circumstances is to be pointless, the relativist must concede that the details of the example could be filled out so that C_{F1} and C_{F2} are *different*.” (Ibid.)

If the circumstances to which the monkey's judgment about the colour of the first fruit are relativized are different from the circumstances to which the monkey's judgment about the colour of the *second* fruit are relativized, then – so the objection goes – the monkey ascribes different colour properties to the two fruits. Accordingly, Byrne & Hilbert draw the following conclusion:

“According to the relativist, the first fruit seemed to you to have a *different* color than the second, and hence the relativist cannot endorse the simple and obvious explanation of your fruit-eating behavior.” (Ibid.)

In response to this objection, it should, first of all, be pointed out that it is not at all clear in how far a perceiver that – like the Old World monkey from the example – is supposed to lack any sort of colour language could plausibly be regarded as making *judgments* about colour. In such a case, it might seem more plausible to provide an explanation that does not at all refer to judgments, but merely to similarities between phenomenal appearances, that is, to the similarity between the phenomenal colour appearance of the second fruit and the memorized colour appearance of the first one. Such an explanation would not only be available to the absolutist, but also to the relativist, because even if relativized colour judgments in the two cases would refer to different colour properties, the similarity between the phenomenal appearances would still be given.

Moreover, even if it is granted that the example calls for an explanation in terms of colour judgments, the objection still fails, for reasons similar to those presented in dealing with the objections regarding disagreement and illusion. As already mentioned, the most plausible version of colour relativism is one according to which the kinds of circumstances colour judgments are relativized to include all those viewing conditions under which – according to presuppositions of the perceiver – the object in question phenomenally appears (to the kind of perceiver the judgment is relativized to) to have the colour actually ascribed to it by the perceiver. Accordingly, in Byrne & Hilbert's example, it is fairly plausible to assume that, if it is granted that the Old World monkey *judges* the second fruit to have a certain colour, this judgment is not merely relativized to the monkey's actual viewing conditions, but also to the kind of conditions under which the monkey perceived the first fruit, because the monkey – conceived of as a proper judge – probably presupposes that the second fruit would phenomenally appear (to the monkey) to be of the same shade of red under the kind of viewing conditions

under which the monkey perceived the first fruit – at least if the explanation of the monkey’s fruit-eating behaviour given in the example is to be plausible. It should therefore be assumed that the same colour property is ascribed to both fruits, and that the relativist can thereby provide the desired explanation.

The example might, of course, also be modified such that the two fruits phenomenally appear slightly different in colour to the monkey (due to slight changes in lighting conditions), and it might be assumed that the monkey nonetheless recognizes them to be of exactly the same shade of crimson. The relativist can explain this kind of recognition by, once again, making the plausible claim that the monkey – conceived of as a proper judge – implicitly assumes that the second fruit, perceived under the conditions under which the first one was perceived, would phenomenally appear to be of a shade of crimson relevantly similar to the colour of the first fruit (or that the *first* fruit, perceived under the conditions under which the second one is perceived, would phenomenally appear to be of a shade of crimson relevantly similar to the colour of the second fruit). That is, the monkey will then compare the colours of the two fruits either (i) by judging both colours relative to the kind of circumstances under which the second fruit is perceived, or (ii) by judging both colours relative to the kind of circumstances under which the first fruit was perceived, or maybe by doing both (i) and (ii). The kind of recognition of similarity mentioned in the description of the example amounts, regarding each of the pairs of judgment described under (i) and (ii), to a similarity of the ascribed phenomenal appearances. The two colour judgments described under (i) are relativized to the same kind of circumstances, and therefore refer to the same colour property; and the same holds regarding the two colour judgments described under (ii). So the relativist can, again, provide the desired explanation.

5.3. Some options for resisting the view

It has occurred in the previous section that the problem of standard variation poses a serious threat to views about colour that appeal to the notions of standard perceivers and standard circumstances – be it theories according to which colours are dispositions to phenomenally appear in a certain way to standard perceivers under standard circumstances, theories according to which colours are the physical bases of these dispositions, etc. It has also occurred that colour relativism appears to provide a

relatively elegant solution to the problem of standard variation, and that a specific version of the view – namely one according to which the relata are determined by the speaker’s (or the interlocutors’) presuppositions and are therefore not restricted to (and do not even have to include) the speaker herself, or the speaker’s circumstances, or circumstances very similar to those the speaker finds herself in – can be convincingly defended against certain objections that have been raised against the view.

However, the variation phenomena discussed in the previous section still leave some room for various versions of colour absolutism. One such version may be introduced by drawing an analogy between the case of colour and, for instance, the case of heat. In the latter case, the claim that the phenomenon in question can be identified with some kind of physical, numerically describable, phenomenon – namely, molecular motion, roughly speaking – is relatively undisputed. As in the case of colour, the phenomenon of heat is associated with certain phenomenal qualities. More exactly, while different colours are associated with certain distinctive phenomenal qualities, the same is true for different degrees of heat, that is, coldness, warmth, hotness, etc. One might therefore ask why the two phenomena should theoretically be treated in fundamentally different ways.

The analogy between the two phenomena becomes even more striking in the light of the fact that all the standard variation phenomena considered above regarding the case of colour (except for Jackson’s magazine example) also occur, in an analogous form, regarding the case of heat. That is, objects with the same intrinsic physical constitution may, if touched by different persons, or by the same person at different times, or by the same person with different parts of the skin, phenomenally appear differently regarding fine-grained degrees of heat, depending, for instance, on the initial temperature of the respective part of the person’s skin.

A first step in vindicating an absolutist position regarding colour on the basis of the analogy between colour and heat could therefore consist in conceding that in the case of heat, ordinary-language terms like ‘cold’ and ‘hot’ may, in the light of the (supposed) respective standard variation phenomena, best be treated in a relativist way. However, this does not at all amount to a full-blown relativization of the phenomenon of heat. Rather, the identification of this phenomenon with certain numerically describable, *non*-relative, properties of objects is compatible with the relativization of ordinary-language terms like ‘cold’ and ‘hot’, because the language of physics provides its own vocabulary

for exact and objective descriptions of degrees of heat – descriptions that ordinary-language terms like ‘cold’ and ‘hot’ are not supposed to provide.

Arguably, in the case of colour, matters are a bit different, due to the fact that the visible colour spectrum does not phenomenally appear as linear as the (middle part of the) spectrum of different degrees of heat. However, this phenomenal disanalogy does by no means constitute a knock-down argument against describing colours in terms of linear, numerically describable properties of objects (see Churchland 2007 for an elaborated attempt). Given the concession that ordinary-language colour terms may well deserve a relativist treatment, one would also have to concede that it may not be possible to provide a *mapping* of the numerically describable properties in question on ordinary-language colour terms. However, given the analogy in question between the case of colour and the case of heat, it does not seem far-fetched to regard colour as a phenomenon the exact and objective measuring of which can only be provided in scientific terms. On this radically physicalist view, phenomenal appearances at most play the role of *indicating* the correct description of an object’s colour. The view might therefore be regarded as a promising absolutist option for escaping the standard variation problem.

However, this option – no matter how promising as an attempt at defending a version of colour absolutism in the face of standard variation phenomena – fails to vindicate PI_{unres} . The reason is that, even though the option in question is incompatible with Rampant Variation given that the term ‘colour’ in the formulation of Rampant Variation is supposed to refer to the meaning of the appropriate scientific vocabulary, it is *not* incompatible with Rampant Variation given that the term ‘colour’ in the formulation of Rampant Variation is supposed to refer to the meaning of ordinary-language colour terms like ‘red’ and ‘crimson’. That is, the radically physicalist option in question is still compatible with the following claim:

(Rampant Variation_{OL}) In the actual world, there is not even a rough distinguishing correlation between the meaning of any ordinary-language colour term and any hidden surface (or light emission) properties of things.

Obviously, this reading of Rampant Variation, and not the more scientific reading mentioned before, is the one relevant regarding PI_{unres} . And since this reading is

compatible with the radically physicalist option in question, the latter does not provide a way to defend $P1_{unres}$ against arguments for colour relativism.

Another option for vindicating a version of colour absolutism in the face of standard variation phenomena is to embrace colour primitivism – the view that colours are intrinsic, qualitative, non-reducible properties of objects (defended, for example, in Campbell 1993 and, more recently, in Gert 2008). But regardless of the prospects of this kind of theory, it is obvious that resorting to this option is of no help for defenders of $P1_{unres}$, since it is, like the radically physicalist option considered above, compatible with Rampant Variation_{OL}.

5.4. Resorting to modified versions of P1

Given the fact that the above-mentioned attempts at resisting colour relativism fail to vindicate $P1_{unres}$, I will now (as already announced in section 5.1) explore the prospects of vindicating the argument from analogy by resorting to different versions of P1 as well as to corresponding modifications of the scenario to which the argument to analogy refers.

A first such option might involve a modification of the scenario of Mary and Norman such that Mary teaches Norman all there is to know about when objects are crimson for certain perceivers under certain circumstances, involving all relevant information about intrinsic microscopic properties of the object in question, viewing circumstances, and kinds of perceivers. Accordingly, P1 would be reformulated as follows:

($P1_{res1}$) In virtue of possessing the reliable skill of determining whether an object is crimson for a certain perceiver under certain viewing conditions – given sufficient information about intrinsic microscopic properties of an object, about the viewing conditions, and about the kind of perceiver in question – Norman understands the word ‘crimson’.

Defenders of this first option for vindicating the argument from analogy would have to admit that the scenario referred to in $P1_{res1}$ is considerably more outlandish than the original version would appear to be (at least from the perspective of an absolutist

position according to which the original version is metaphysically possible). That is, not only does the kind of knowledge in question – due to the supposedly required degree of complexity – appear to be hardly acquirable by any human being; but it would also require Mary and Norman to be in some God-like epistemic position regarding phenomenal states of different kinds of perceivers – an epistemic position that might, for instance, be taken to involve knowledge about the possibility of, and the conditions for, spectral inversion.³⁷ I take it, therefore, that resorting to $P1_{res1}$ and the corresponding modification of the scenario is not an attractive option for defending the argument from analogy.

A second option might involve a modification of the scenario such that Mary does not teach Norman anything about microphysical surface structures, wavelength, etc., but instead about the kind of activation of cone cells that causes sensations of crimson for human perceivers. Like other so-called trichromats, humans possess three types of cone cells in the retina, and the ratio of activation of these cells completely determines (at least in the case of well-functioning neural systems) the respective hue sensation (see Cornsweet 1970, pp. 243ff., for a detailed description). Modifying the scenario in this way yields the following modification of P1:

($P1_{res2}$) In virtue of knowing about the ratio of activation of cone cells that causes sensations of crimson for human perceivers with well-functioning neural systems, Norman understands the word ‘crimson’.

Obviously, the kind of knowledge mentioned here is much more easily acquirable than the kind of knowledge described in $P1_{res1}$ – and it has supposedly actually been acquired by scientists. However, given the relatively sparse character of this kind of knowledge – due to the restriction to human perceivers with well-functioning neural systems – $P1_{res2}$ requires conditions for linguistic understanding considerably more permissive than those required by $P1_{unres}$, and it might well be rejected on these grounds. (And the same holds, to an even larger extent, regarding a possible version of $P1_{res2}$ that does not concern human perceivers in general, but merely Norman himself, or kinds of

³⁷ One might try to avoid these epistemic worries by endorsing a functionalist view regarding phenomenal appearances. I simply grant to opponents of the argument from analogy that this strategy would be of no help in trying to vindicate $P1_{res1}$ (or any of the versions of P1 to be considered below). In the next section, however, I will argue that endorsing a functionalist view regarding phenomenal appearances would be of no help for colour relativists in dealing with the problems that befall their own view in the light of standard variation phenomena.

perceivers relevantly similar to Norman; the corresponding knowledge – so one might reasonably object – could hardly suffice for the linguistic understanding in question.) Of course, one might resort to further modifications, involving knowledge that does, for instance, not only concern human perceivers, or trichromats in general, but also tetrachromats, or even further possible kinds of perceivers. But as in the case of $P1_{res1}$, this would require Mary and Norman to be in some God-like epistemic position regarding phenomenal states of different kinds of perceivers. (How is it, for instance, even in principle possible for humans to come to know whether pigeons, and other tetrachromats, ever have phenomenal colour sensations similar to that of trichromats?) Moreover, $P1_{res2}$ itself might already be taken to require Norman to be in some God-like epistemic position – namely, again, regarding the possibility of, and the conditions for, spectral inversion. (This point does not hold regarding a possible version of $P1_{res2}$ that merely concerns Norman himself, or kinds of perceivers relevantly similar to Norman. But as already mentioned above, the kind of knowledge described in such a version of $P1_{res2}$, or of $P1_{res1}$, would hardly be general enough to suffice for the linguistic understanding in question.) I take it, therefore, that this second option for defending the argument from analogy does not fare significantly better than the first one.

Another option might involve a modification of the scenario such that Mary teaches Norman about a trans-species kind of brain state which corresponds, maybe as a matter of metaphysical necessity, to occurrences of all and only those phenomenal states that involve a sensation of crimson. Obviously, though, this option does not only head into even more speculative terrain than the other two considered above, it is also subject to the same kind of epistemic worry as most of the other versions already considered.

It should be concluded that resorting to modified versions of P1 is not a promising strategy for vindicating the argument from analogy in the face of the considerations in favour of colour relativism presented above. However, in the following section, I will argue that the occurrence of standard variation phenomena does not only pose a threat to any absolutist view appealing to standard perceivers and standard circumstances; rather, their occurrence involves an equally serious problem for colour relativism.

5.5. How standard variation phenomena turn against colour relativism

As already described, colour relativism seems to provide an elegant solution to the challenge posed by the occurrence of standard variation phenomena. The reason is that while the most commonly endorsed rival theories – according to which colours are dispositions to phenomenally appear in certain ways to standard perceivers under standard circumstances, or the bases of these dispositions, or corresponding second order properties – all appeal to the notions of standard perceivers and standard circumstances, colour relativists seem to offer an account that can dispense with these notions, without thereby being committed to colour scepticism or colour primitivism. However, there is an important feature that colour relativism shares with most of its competitors, including those commonly endorsed rival theories mentioned above. This feature consists in an appeal to certain phenomenal appearances. That is, what makes it, according to colour relativism, the case that an object has a certain colour for certain perceivers under certain circumstances is that the object in question *phenomenally appears* in a certain way to these perceivers under these circumstances.

Of course, alternative versions of the view are conceivable. The most obvious alternative would be a version of the view according to which what makes it the case that an object has a certain colour for certain perceivers under certain circumstances is not that the object in question phenomenally appears in a certain way to these perceivers under these circumstances, but rather that the perceivers in question (implicitly) *judge* the object to be of a certain colour under the circumstances in question. Such a version of the view, however, involves quite unwelcome commitments. Imagine, for instance, a person looking at an object under dim moonlight. Due to the sparse lighting conditions, the object phenomenally appears dark grey to her. She judges the object to be red, since her memory tells her that it appeared red to her under better lighting conditions. Imagine further that her memory misleads her in this case, and that the object in fact appeared green to her under better lighting conditions. According to the alternative version of colour relativism described above, it would be true to say that the object in question is red, given that such an ascription is relativized to the person and the circumstances in question. (Of course, it would also be true to say that the object is green, if this ascription is relativized to better lighting conditions.) I take this result to be quite unwelcome. According to the standard version of colour relativism, ascriptions of redness to the object might also be true, namely, if these ascriptions are relativized to

unusual lighting conditions, or unusual perceivers, of a certain kind. However, this commitment seems clearly more acceptable than the one proponents of the alternative version are faced with, since the person's judgment that the object is red in the example described above is simply due to an error, and should therefore not be regarded as providing a basis for a true colour ascription, even if this ascription is merely relativized to the peculiar circumstances described in the example.

The alternative version in question faces an additional problem, which I take to be even more serious than the one just described. Given the general relativist claim that colour ascriptions are relativized to kinds of perceivers and circumstances, it would be incoherent to deny that the person's redness judgment in the above example is also relativized in a certain way. But then, since the notion of judgment was supposed to figure in an explanation of what makes it the case that an object has a certain colour for certain perceivers under certain circumstances, and since the relevant kind of judgment is itself relativized to kinds of perceivers and kinds of circumstances, then proponents of the view in question – in order to be coherent – would have to claim that the relevant judgment is itself a judgment about which colour certain perceivers would *judge* an object to have under certain circumstances, and this latter judgment – also relativized to kinds of perceivers and kinds of circumstances – would itself have to be regarded as a judgment about which colour certain perceivers would judge an object to have under certain circumstances, and so on. Thus, the supposed explanation of what makes it the case that an object has a certain colour for certain perceivers under certain circumstances appeals to a notion – namely, the notion of colour judgment – which is itself in need of the same kind of explanation it was supposed to provide. Proponents of the alternative version in question are thereby faced with a vicious regress. For this reason, taken together with the objection presented in the previous paragraph, the view should be rejected.

Further alternative versions of colour relativism might be proposed – versions according to which what makes it the case that an object has a certain colour for certain perceivers under certain circumstances is neither the object's phenomenal appearance to these perceivers under these circumstances nor the perceiver's colour judgment regarding the object under these circumstances. But I take it that no such version suggests itself as promising. Therefore, the burden of proof lies with those who claim that a plausible version of the kind just described is available. And since versions appealing to the perceivers' colour judgments have turned out implausible, it seems that

the plausibility of colour relativism amounts to the plausibility of its standard version – the version according to which an object has a certain colour for certain perceivers under certain circumstances in virtue of the object’s phenomenal appearance to these perceivers under these circumstances. I will now argue that this latter version of the view faces problems analogous to those presented by colour relativists against standard rival theories. That is, I will argue that colour relativism is, in the end, not better off in the face of standard variation phenomena than standard versions of colour absolutism.

My point against versions of colour relativism that appeal to an object’s phenomenal appearance is the following. When presented with a relativist account of the meanings of colour terms, according to which the correct use of these terms depends on the object’s phenomenal appearance to those perceivers and circumstances the colour ascription is relativized to, someone who is about to learn a new colour term ‘c’ will wonder: *which* kind of phenomenal appearance to certain perceivers under certain circumstances does an object have to possess in order to have colour c for these perceivers under these circumstances? And one might further ask: how did it come about that one particular, and not some other, kind of phenomenal appearance is involved in the meaning of ‘c’? I take it that for an account of colour terms to be tenable, it must not commit one to the claim that these questions constitute unsolvable puzzles. It should be remembered that, apart from Jackson’s example of the magazine patches – which turned out to be a surmountable threat for standard versions of colour absolutism –, all the standard variation phenomena considered so far do not concern coarse-grained colours like red, blue, orange, etc., but merely relatively specific colour *shades*. And I now claim that regarding relatively specific colour shades, the questions just mentioned *do* (in the light of standard variation phenomena) constitute unsolvable puzzles, and thereby reveal a serious shortcoming of relativist accounts of the meanings of colour terms.

To illustrate the problem the relativist is faced with, consider, as an example of a term for a specific colour, ‘true blue’, as coined by Tye 2006a. This term is supposed to refer to a shade of blue not at all tinged with yellow or red. Considered as a term for a fully saturated colour, ‘true blue’ might thereby be regarded as a term for a maximally specific colour shade. According to colour relativism, an object is true blue for certain perceivers under certain circumstances in virtue of phenomenally appearing in a certain way to these perceivers under these circumstances. Now, consider again the two questions posed above. Given the occurrence of all the different kinds of variation

phenomena described in section 5.2, it is plausible to assume that there is no single colour shade commonly associated with the term ‘true blue’ by most competent language users. That is, given the high sensitivity to variations of appearance regarding specific colour shades, it appears quite improbable that two given language users have come to associate the term ‘true blue’ with the same phenomenal appearance. But if a relativist account of this kind of term is correct, then there is one, maximally specific, phenomenal colour appearance in virtue of which an object counts as true blue for certain perceivers under certain circumstances. And in order not to treat the two questions posed above as unsolvable puzzles, the relativist seems to have no choice but to resort to the notion of standardness in some way or other. That is, in order to provide an answer to the question *which* phenomenal appearance an object must have in order to count as true blue for certain perceivers under certain circumstances, and to the question of how it came about that one particular, maximally specific, colour shade, and not some other, is involved in the meaning of ‘true blue’, it seems inevitable to point to certain standard viewing conditions that reveal an object’s real colour to certain standard kinds of perceivers. (And even though these two kinds of questions appear more pressing in the case of relatively specific colours, they can of course also be asked regarding more general colour terms like ‘red’, ‘blue’, etc., and it seems that in these cases, relativists would be forced to provide the same kind of answer as the one just mentioned.)

At this point, the relativist might try to resort to a version of the view according to which the meanings of terms like ‘true blue’ do not involve one single kind of phenomenal appearance for all perceivers. Rather, one might hold, the kind of phenomenal appearance in virtue of which it is correct to ascribe a certain colour to an object for certain perceivers under certain circumstances may vary from one speaker to another (or, alternatively, from one perceiver to another). However, many proponents of this kind of view would presumably not be willing to deny that some ways of associating a colour term with a certain kind of phenomenal appearance are deviant, and that the respective person should, *ceteris paribus*, not be regarded as understanding the term. And when it comes to determining which ways of associating a given colour term with a kind of phenomenal appearance count as deviant, it is – once again – not at all clear how an appeal to the notions of standard perceivers and standard circumstances can be avoided.

Alternatively, it might be held that the meanings of terms like ‘true blue’ do not involve appearances with a particular phenomenal character, but instead appearances realizing a certain functional role – a role that might in principle be realized by appearances with any phenomenal character. (A functionalist view of this kind, within an absolutist rather than a relativist framework, is, for instance, proposed by Glüer 2007.³⁸) Given such a conception of the meanings of terms like ‘true blue’, the relativist would claim that an object is true blue for certain perceivers under certain circumstances iff the object’s phenomenal appearance to these perceivers under these circumstances realizes a certain functional role. This way of conceiving of the meanings of colour terms has, for instance, the consequence that spectral inversion has no influence on which colour an object has for certain perceivers under certain circumstances, as long as the inversion has no impact on the functional roles realized. The most self-suggesting candidate for a criterion for realizing the relevant kind of functional role involves what Glüer 2007 calls “color classifying behavior” (p. 123). That is, the most self-suggesting way for the relativist to implement a functionalist account of the meanings of colour terms would be to hold that an object has colour *c* for certain English-speaking perceivers under certain circumstances iff the colour appearance caused by the object under these circumstances is such that – as Glüer also puts it, “in the absence of beliefs to the effect that experiences are not to be trusted” (p. 124) – these perceivers hold ‘the object has colour *c*’ true.

On the face of it, it might seem that this version of colour relativism avoids any appeal to the notions of standard perceivers and standard circumstances. But this verdict would be too hasty, for the following reason. Obviously, speakers of English differ with respect to the exact ways in which they learn certain colour terms. And speakers may differ in their colour ascriptions merely due to differences in how they acquired the respective colour terms. Now, it is to be expected that relativists would not be willing to

³⁸ More precisely, Glüer proposes a functionalist account of the meanings of colour *sensation* terms, which are to be distinguished from colour terms, and which, in turn, are often held to be required for defining the meanings of colour terms without running into circularities (see especially Peacocke 1983, ch. 2, and 1984). Glüer proposes a functionalist account of the meanings of colour sensation terms as a solution to certain further circularity worries. I will not take a stance regarding whether the resulting account of the meanings of colour terms should be regarded as the most promising version of an absolutist account of the meanings of colour terms. As I will argue below, a relativist account of the meanings of colour terms incorporating a functionalist account of the meanings of colour sensation terms cannot avoid the problems regarding standard variation phenomena discussed above. And a reason analogous to the one to be presented below for the claim just stated can be provided for the claim that absolutist accounts of the meanings of colour terms incorporating a functionalist account of the meanings of colour sensation terms can also not avoid the problems in question.

hold that the object has colour *c* under certain circumstances for *any* English-speaking perceiver whose colour experience realizes the functional role described above, completely regardless of how this perceiver learned to use the respective colour term. Rather, it would be more plausible to allow for some ways of learning to use colour terms to be inadmissibly deviant, such that, given the resulting deviant usage of the term, having a colour experience realizing the functional role in question is not sufficient for the object having colour *c* for the perceiver in question under the circumstances in question – simply because the perceiver in question does not count as a competent user of the term. In order to incorporate this kind of exception, without appealing to the phenomenal character of the relevant experience, it seems that the functionalist account in question would, once again, have to appeal to the notion of standardness, namely, to the notion of a standard user of the respective colour term – that is, either to the notion of a speaker who has learned the term under certain standard circumstances, or simply to the notion of a speaker who is disposed to agree to certain colour ascription sentences under standard circumstances.

Unfortunately for proponents of colour relativism, resorting to the notion of standardness in the ways described brings about problems similar to those that commonly endorsed versions of absolutism are faced with in the light of standard variation phenomena. (Regarding the functionalist account just described, resorting to the notion of standard circumstances would be problematic in so far as, in the case of specific colour shades like true blue, colour ascriptions are highly sensitive to changes in viewing conditions even within the range of conditions that would seem to deserve the label ‘standard’.³⁹) And since it is not at all clear how colour relativists could plausibly avoid resorting to the notion of standardness, it should be concluded that they are, after all, not significantly better off in the face of standard variation phenomena than their absolutist opponents.

The same problem besets certain other theories of colour that initially seemed not to be affected by the occurrence of standard variation phenomena. That is, in order for colour primitivism – the view that colours are irreducible, qualitative properties of objects – not to leave it completely mysterious *which* phenomenal appearance an object must have in order to count as true blue for certain perceivers under certain

³⁹ As already alluded to in fn. 38 above, an analogous problem also occurs with respect to *absolutist* accounts of the meanings of colour terms involving a functionalist account of the meanings of colour sensation terms.

circumstances, and how it came about that one particular, maximally specific, colour shade, and not some other, is involved in the meaning of ‘true blue’, an appeal to the notion of standardness seems inevitable. And the same holds for any kind of projectivist account according to which even though there are phenomenal colour concepts, no object falls under them.

Thus, the problem of standard variation is much more widespread than it may initially appear. Since it besets colour relativism to the same extent as its standard competitors, the fact that – as described in section 5.1 – colour relativism is in tension with P1 should no longer be particularly worrying for proponents of the argument from analogy. However, the question of how to deal with standard variation phenomena still remains. In the following section, I will describe the kinds of theories that can accommodate these phenomena, as well as the relation of these theories to the argument from analogy.

5.6. The remaining options

It has turned out that the problem of standard variation poses a serious threat to all theories of the meanings of specific colour terms like ‘true blue’ that either directly appeal to the notion of standardness or involve the claim that the meanings of colour terms like ‘true blue’ involve specific phenomenal appearances. Fortunately for proponents of the argument from analogy, the only promising version of colour relativism falls under the latter kind of theory, and is thereby not better off in the light of standard variation phenomena than its most commonly endorsed competitors. However, it remains to be examined which theories of the meanings of colour terms can still be regarded as tenable in the light of the phenomena in question, and how these theories relate to the argument from analogy.

The most obvious possibility to accommodate the phenomena in question is to endorse a radically eliminativist position regarding the meanings of colour terms. In contrast to projectivist accounts – according to which colour terms are meaningful, but no object has any colour – a radically eliminativist account involves the claim that colour terms are meaningless. Of course, given the initial counter-intuitive character of such an account, it should merely be regarded as a last resort, to be endorsed only if all alternative views turn out untenable.

Assuming that the radically eliminativist position just described is correct, it would have to be concluded that the argument from analogy is unsound – for if the terms ‘crimson’ and ‘red’ are meaningless, premise 1 would have to be regarded as false, since the sentence ‘all crimson things are red’ would not express anything that can be known or understood. Nevertheless, as stated above in the introduction, my aim in putting forward the argument from analogy is to provide an argument for the claim that colour inclusion sentences do not express a priori truths. And this claim would obviously be vindicated within a radically eliminativist picture, for within this picture, colour inclusion sentences do not express truths at all. Thus, this way of accommodating standard variation phenomena would commit one to accepting the truth of my central claim, even though for a rather surprising reason.

However, independently of the purpose of putting forward the argument from analogy, its proponents still have reason to be optimistic. For another way of accommodating standard variation phenomena is available – one that is not only compatible with my central thesis, but also with the soundness of the argument from analogy. As already mentioned, all the standard variation phenomena considered above, with the exception of Jackson’s example of the magazine patches, do not concern coarse-grained colours like red, blue, orange, etc., but merely relatively specific colour *shades*. And since, as described in section 5.2, the example of the magazine patches may well be regarded as a surmountable threat, a possible way to accommodate the phenomena considered so far is to opt for a theory that treats the meanings of terms like ‘true blue’ differently from the meanings of terms like ‘red’ and ‘green’. That is, one might opt for a theory according to which relatively general colour terms refer to colours possessed by objects in an absolute sense, but according to which relatively *specific* colour terms are meaningless. (One may, of course, try to stipulate that a certain specific colour term refer to a particular colour shade experienced at a given moment; but defenders of the view just described could simply hold that the resulting artificial term could not be a colour term, but merely a term referring to a certain kind of appearance.)

Such a fundamentally heterogeneous treatment of terms like ‘red’ and ‘green’, on the one hand, and terms like ‘true blue’, on the other, may initially appear quite unattractive, since the heterogeneity in question constitutes a considerable theoretical cost. However, after considering all the variation phenomena described above, a different perspective on the heterogeneous treatment in question arises. That is, the fact

that some rampant variation of phenomenal appearance occurs at the level of specific colour shades, but – at least according to what has been considered so far – not at the level of coarse-grained colours, might well be taken to justify the heterogeneous treatment in question.

A possible worry regarding this view concerns the location of the threshold between the level of generality at which colour terms refer to colours of objects in an absolute sense and the level of specificity at which colour terms are meaningless. Should it turn out that any possible way of locating this threshold involves an arbitrary choice, then the view has to be regarded as facing a serious problem. However, given the empirically based motivation for endorsing the view, it seems plausible that the most adequate way to deal with this issue is to hold that the location of the threshold is to be determined empirically, that is, by observing at which level of generality variation phenomena make choices for standard kinds of perceivers and standard kinds of circumstances inadmissibly arbitrary. For proponents of the argument from analogy, it will be particularly interesting whether the term ‘crimson’ is general enough in order not to fall below the threshold. If it does fall below the threshold, proponents of the heterogeneous account would have to regard the argument from analogy as unsound, for reasons already mentioned. However, they might still regard the argument as valid when applied to colour inclusion statements that only involve colour terms of a sufficiently high level of generality. And, as stated above, even though the argument from analogy were unsound when applied to colour terms of an insufficiently high level of generality, these cases would not constitute counterexamples to my central thesis (since the corresponding sentences would not express truths, and would thereby not express a priori truths). But proponents of the argument from analogy may be optimistic, for the variation phenomena considered so far do not constitute a good reason for denying that some reasonable choice of standard kinds of perceivers and circumstances is available regarding the term ‘crimson’.

Another worry could arise from the suspicion that, in contrast to what has been concluded so far, standard variation phenomena also occur regarding general colours like blue and red. Cohen et al. 2006 – in response to a proposed solution to the problem of standard variation regarding specific colour shades presented in Tye 2006a⁴⁰ – point

⁴⁰ In response to the problem of standard variation, Tye proposes an absolutist solution according to which for each colour shade, there can in principle be colour experiences that accurately represent the shade. Tye’s explanation of the problem consists in the claim that for evolutionary reasons, humans are

to a study by Malkoc et al. 2005, which supposedly shows that the kind of standard variation phenomena that undoubtedly occur regarding specific colour shades also occur, to some significant extent, at the level of coarse-grained colours:

“[T]here is in fact variation in colour experience among standard perceivers in standard circumstances even for colours at the level of grain of blue, purple, orange, and the like.” (Cohen et al. 2006, p. 339)

Cohen et al. go on to give the following report of the study in question:

“[I]n a recent colour-naming study with 34 hue samples and eight colour categories, seven hue samples were labelled blue by at least one subject, but there was 80% consensus on only two (Malkoc et al. 2005). All seven samples were assigned to coarse-grained colours other than blue at least once. Examination of figure 4 of Malkoc et al. 2005 shows that [t]his degree of overlap was characteristic of all of the coarse-grained categories.” (Cohen et al. 2006, p. 339)

In order to adequately evaluate in how far these reported results constitute a problem for an absolutist view regarding coarse-grained colours, it is useful to have a closer look at Malkoc et al.’s own report of the study. The mere fact that there was a low overall degree of consensus regarding the categorization of the colour samples does not necessarily mean that the problem of standard variation also occurs regarding coarse-grained colours, since it might, for example, be the case that the colour samples were chosen randomly and that all the chosen samples happen to be such that none of the participants of the study regarded any of them as a *clear* sample of one of the eight colours in question. That is, all of the chosen samples might happen to be borderline cases. A closer look at the study reveals that the choice of the 24 (not 34, as reported by Cohen et al.) samples resulted from a division of the visible colour spectrum into equal intervals (cf. Malkoc et al. 2005, fig. 4, p. 2163). This might seem to prevent the possibility of the samples all happening to be borderline cases. However, given that the intervals range over the whole visible spectrum, the distance regarding phenomenal

not designed to accurately represent specific colour shades. Accordingly, he laxly claims that “God knows precisely which hue chip 527 has, but we may very well never know” (Tye 2006a, p. 177). It should be noted that Tye’s short description of the proposal leaves it quite obscure in virtue of which kind of fact phenomenal appearances correctly represent specific colour shades. He would presumably have to claim that phenomenal appearances correctly represent specific colour shades in virtue of some possible evolutionary extension of our colour detection faculties. But it is not at all clear whether among the many possible ways in which such an evolutionary process could occur (where different possible ways would result in different phenomenal representations of the same colour samples), *one* could reasonably be singled out as privileged, such that it is *this* possible extension in virtue of which specific colour shades can already right now in principle accurately represent specific colour shades. In a recent book symposium on Cohen 2009, Tye offers some further brief remarks in favour of the epistemicist position just described (see Tye 2012, pp. 298-300). See Cohen 2012, pp. 313-316, for a refutation of Tye’s remarks.

colour appearance between each of the 24 steps is relatively big, and it might therefore still be that it happened to be the case that none of the samples was such that any of the participants of the study regarded it as a *clear* sample of the colour they attributed to it. Malkoc et al.'s report of the study does not provide further information regarding this matter. As indicated by Cohen et al., none of the 24 samples was attributed the same colour by all participants. (59 students participated in this study; cf. Malkoc et al., p. 2156). Moreover, Malkoc et al. categorize all of them as "color-normal" (ibid.), that is, as what would normally be regarded as standard colour perceivers, and all participants adapted for one minute to the grey background against which the samples were shown (cf. ibid.). The lack of complete consensus concerning any of the 24 samples might therefore well be regarded as a striking result, calling into question the tenability of the heterogeneous account.

It should be noted that the result of the study does not prove the occurrence of significant differences among the ways in which the hues phenomenally appeared to the participants. Rather, the study merely shows differences among the ways in which the participants *labelled* the hues. These differences might have arisen from different ways in which the participants have learned to associate the respective colour terms with kinds of phenomenal appearances, and are therefore still compatible with the assumption that there were no significant differences among the ways in which the hues phenomenally appeared to the participants. However, even if the result is explained merely on the basis of differences among ways of acquiring colour terms, it should still be regarded as posing a threat to the restricted absolutist position in question, because it might seem that proponents of this position would then be pressed to characterize, in a non-arbitrary way, some of the participants as standard regarding their ways of acquiring the colour terms in question, and some others of them as non-standard.

However, regardless of whether the result of the study is explained on the basis of differences among phenomenal appearances or on the basis of differences among ways of associating colour terms with certain phenomenal appearances, it is still compatible with the plausible assumption that, if any of the participants had generated a sample that she herself regards as a maximally clear sample of any of the eight colours in question (except, maybe, yellow-green – the only of the four binary hues for which there is no common ordinary-language term in English), all the other participants would have attributed the same of the eight colours to this sample. One might even claim that the same of the eight colours would have been attributed by all the participants under all

kinds of viewing conditions that could reasonably be defined as standard (even though this claim is of course much more vulnerable than the one mentioned before). The absolutist regarding coarse-grained colours might therefore defend her view against the results of Malkoc et al.'s study by conceding that coarse-grained colours might be vague to a higher degree than expected, but by nevertheless insisting that there are core samples determinately possessing particular coarse-grained colours in an absolute sense.⁴¹

These remarks in defense of the heterogeneous account do, of course, not yet constitute a convincing case in favour of it, and I do not aim at providing such a case. Instead, my aim is merely to point out that the account is among the not too far-fetched options for accommodating standard variation phenomena, and that it is either compatible with the argument from analogy or (like the radically eliminativist option considered above) incompatible with the falsity of my central thesis.

Those who are willing to avoid the threat posed by the occurrence of standard variation phenomena, but unwilling to endorse one of the two options already considered, may try to argue for another account – one according to which objects possess both coarse-grained and fine-grained colours, but according to which (i) no appeal to the notions of standard perceivers and standard circumstances are needed to spell out the meanings of colour terms, and (ii) the meanings of colour terms do not involve certain phenomenal appearances. The view I have in mind here is based on the claim that colour terms are in a certain way analogous to proper names. That is, analogously to the currently most commonly held position regarding proper names, one might hold that the fact that colour terms refer to certain properties of objects is due to a causal chain that leads from some past act (or group of acts) of baptism to current uses of the respective term, and that this causal chain is held together by speakers' intentions to use the respective term in the same way as those they have picked it up from.

Of course, this account of the meanings of colour terms involves more complications than the analogous account of proper names. In the case of colour terms, it is, for instance, much less obvious how to locate uses of the respective term that may count as fixing the term's reference. And it is not at all clear whether any promising

⁴¹ In response to Cohen et al.'s objection, Tye tries to reconcile the possible assumption that each of the colour samples falls into exactly one of the eight colour categories with the possibility that none of the participants of the study misrepresents any of the samples. He does so by appealing to the notion of *under-*, instead of *mis-*, representation (cf. Tye 2006b, p. 343). See Cohen et al. 2007, pp. 164f., for a refutation of this attempt.

recipe for locating the reference-fixing uses could dispense with the notions of standard perceivers and standard circumstances. If it could not, then endorsing the view would be no option for avoiding the threat posed by the occurrence of standard variation phenomena. However, regardless of whether the view can dispense with the notion of standardness, and whether the view turns out plausible, it is certainly not in tension with the argument from analogy.

Thus, all options considered so far for avoiding the threat posed by the occurrence of standard variation phenomena are either compatible with the argument from analogy or incompatible with the falsity of my central thesis. I take it that the burden of proof now lies with those who claim that any other at least remotely promising option for avoiding the threat posed by the occurrence of standard variation phenomena is in tension with the argument from analogy.

5.7. Summary

In section 5.1, I have argued that a successful case for colour relativism would undermine a certain version of P1, which I have called 'P1_{unres}', and that the occurrence of standard variation phenomena therefore threatens the argument from analogy, as far as these phenomena favour colour relativism over colour absolutism. In section 5.2, I have argued that the phenomena in question indeed provide a strong basis for rejecting the most commonly endorsed versions of colour absolutism, and that a certain version of colour relativism appears to provide an elegant solution to the challenge posed by the phenomena. In section 5.3, I have considered several non-relativist options for accommodating the phenomena and argued that these options – even if successful in this respect – would not vindicate P1_{unres}. In section 5.4, I have argued that there is no promising way of resorting to different versions of P1 in order to vindicate the argument from analogy. In section 5.5, I have argued that, even though standard variation phenomena have often been taken to support colour relativism, the phenomena in question are not less threatening for proponents of a relativist treatment of colour terms than for proponents of the most commonly endorsed versions of colour absolutism – an observation that neutralizes the threat posed by the occurrence of the phenomena for proponents of the argument from analogy. Finally, in section 5.6, I have argued that the remaining options for avoiding the challenge posed by the phenomena are either

compatible with the argument from analogy or incompatible with the falsity of my central thesis that colour inclusion sentences do not express a priori truths.

This completes my examination of the prospects of the argument from analogy, as well as my case for the claim that colour inclusion sentences do not express a priori truths. Having defended the argument from analogy against various objections in chapters 2-4, the considerations presented in the previous sections now allow me to conclude that my central thesis has turned out true, since the argument from analogy – whose soundness, together with the assumptions formulated in the introduction, entails the truth of my central claim – is either sound, or unsound for reasons incompatible with the falsity of my central thesis.

Conclusions

As announced in section 0.1, my central aim in the dissertation has been to argue that what I have called colour inclusion sentences (that is, non-logically-true sentences expressing a relation of inclusion, or of exclusion, of colour properties) do – contrary to what is widely held among defenders of the a priori – not express a priori truths. My central argument for this claim, which I have called the argument from analogy, has been formulated and tentatively motivated towards the end of chapter 1. Even though the conclusion of this argument merely refers to one particular example of a colour inclusion sentence, the argument can easily be applied, with equal force, to any other truth expressed by a colour inclusion sentence. And even though the conclusion of the argument from analogy merely concerns one particular way of coming to know a truth expressed by a colour inclusion sentence, significant parts of my argumentation in favour of premise 3 – especially my considerations in section 4.2 – can be generalized to any other way of coming to know truths expressed by colour inclusion sentences.

Those unimpressed by my considerations in section 4.2 might still be tempted to hold that, even though Norman's way of coming to know that all crimson things are red in the original version of the scenario (that is, via an offline judgment involving one's generic memory of the superficial looks of crimson and of red things) is not a priori, the sentence 'All crimson things are red' nevertheless expresses an a priori truth, since – so one might hold – other ways of coming to know that all crimson things are red still have good prospects for deserving the label 'a priori'. The most obvious such way would be that of someone who, instead of having learned the word 'crimson' by being presented to the superficial looks of crimson things, has learned the word by being given a description of the colour's looks involving the term 'red', and who comes to know that all crimson things are red simply by inferring it from this description. However, this strategy for arguing that the sentence 'All crimson things are red' expresses an a priori truth overgeneralizes, since it could as well be used for arguing that the sentence 'Water is H₂O' expresses an a priori truth, given that one can come to understand the term 'H₂O' by being told that it is the chemical structure of water (a structure – so one might imagine – that one already has some knowledge about, just not in connection with the label 'H₂O'), which would, analogously to the above-mentioned case, enable one to come to know that water is H₂O simply by inferring it from the description in question.

Furthermore, it is not at all clear how the argumentative strategy in question could be modified without overgeneralizing in an analogous way. Thus, one may plausibly hold that the argument from analogy does not only constitute a case for the claim that Norman's way of coming to know that all crimson things are red in the original version of the scenario is not a priori, but also for the more general claim that colour inclusion sentences do not express a priori truths.

The argument from analogy has the following basic structure. Its first half (the argument to C1) constitutes a case for the claim that (put in general terms) what would usually be regarded as a paradigm way of acquiring a priori knowledge of truths expressed by colour inclusion sentences (namely, making an offline judgment involving one's generic memory of the superficial looks of the respective colours) sometimes yields knowledge of a truth expressed by a colour inclusion sentence in which experience plays an evidential role. This first conclusion is based on two premises. According to the first of these two premises (again, put in general terms), one might come to understand colour terms (or more precisely: one might come to understand or grasp the entities contributed by these terms to the corresponding objects of knowledge) in a certain deviant way, without associating the respective term with any particular kind of superficial looks, such that, understanding one of the two terms figuring in a colour inclusion sentence in this deviant way, while understanding the other colour term in the usual way, one is still not in a position to know the respective truth. And according to the second premise, the kind of experience missing in order for the subject to be in a position to know the respective truth plays an evidential role in the respective instance of knowledge, since it does not provide for the subject's understanding of the relevant terms (nor for offline reasoning skills corresponding to the subject's way of understanding these terms), which the subject, given the truth of the first premise, already possesses before having the missing kind of experience. The second half of the argument from analogy merely contains one further premise, according to which (again, put in general terms), if experience plays an evidential role in the kind of case in question, it also plays this role in any other instance of knowledge of the relevant kind of truth acquired in the same way (that is, by making an offline judgment involving one's generic memory of the superficial looks of the respective colours).

After introducing the argument from analogy at the end of chapter 1, I devoted the remaining chapters to a defense of it. In chapter 2, I elaborated a possible objection against premise 2 according to which the missing experience in the deviant kind of case

– though not providing for the subject’s understanding simpliciter of the relevant term – still improves or enriches the subject’s understanding of the term, and according to which the missing experience therefore plays a merely enabling role. In response, I argued that the argumentative strategy in question, if it is not to overgeneralize, has to presuppose the truth of premise 1. Accordingly, in chapter 3, I went on to examine the prospects of this premise. In particular, I elaborated a possible objection against premise 1 according to which the primary objects of a priori truth (and thereby the primary objects of truth) are of a certain relatively fine-grained kind, such that what defenders of the apriority of the relevant kind of truth intend to express by colour inclusion sentences is not grasped by the subject in the deviant kind of case. In response, I argued that the objects of truth are of a more coarse-grained kind – a response which also involves a positive case for premise 1. I went on to examine the second half of the argument from analogy in chapter 4, where I provided a positive case for premise 3. Finally, in chapter 5, I questioned a presupposition made throughout the foregoing chapters whose falsity would undermine the truth of premise 1 – the presupposition that there are certain microphysical properties shared by all and only those things to which a certain colour can veridically be ascribed. I argued that all promising options for accommodating the motivation for doubting this presupposition – namely, for accommodating the occurrence of so-called standard variation phenomena – are in fact either compatible with the argument from analogy or incompatible with the falsity of my central thesis that colour inclusion sentences do not express a priori truths. In sum, I argued that the argument from analogy is either sound, or unsound for reasons incompatible with the falsity of the claim which the argument is supposed to support.

To summarize my argumentative strategy as briefly as possible: I argued against the apriority of truths expressed by colour inclusion sentences by showing that understanding or grasp of these truths (together with good corresponding offline reasoning skills) does not guarantee being in a position to acquire knowledge of them, and that some experience in instances of knowledge of these truths therefore always plays an evidential role.

I close by making a few additional remarks on the result of the foregoing chapters. First, it has to be admitted that, while I provided positive cases for premises 1 and 3, I did not provide a positive case for premise 2. That is, as mentioned above, I merely defended this premise against one particular objection. It should, however, also be noted that premise 2 is not only (probably) less controversial than premises 1 and 3, but also

more definitional, and thereby less substantial. That is, in contrast to premises 1 and 3, premise 2 may well be regarded as providing (part of) a definition of one of the central notions, namely, the notion of evidentiality. And I take the definitional assumption made by endorsing premise 2 to be widely accepted, and therefore not to be in need of a positive defense. Moreover, one may even hold that a case for my central thesis can already be made on the basis of premise 1 alone, without endorsing premise 2 or premise 3. That is, the fact that colour inclusion sentences (or the truths expressed by them) can be understood (or grasped) without being in a position to know the respective truth (despite possessing good corresponding offline reasoning skills) – a fact that can easily be argued for on the basis of premise 1 alone – may well be regarded as incompatible with the apriority of this kind of truth, independently of any employment of the notion of evidentiality. At least, having to concede that one can understand (or grasp) colour inclusion sentences (or the truths expressed by them) and have good corresponding offline reasoning skills without being in a position to know the respective truth would certainly be an unwelcome commitment for defenders of the apriority of this kind of truth. Rejecting premise 2 alone (especially on the basis of an unusual notion of evidentiality) should therefore not be regarded as a satisfactory way to defend the apriority of the kind of truth in question.

Second, one might suspect that the truth of my claim that experience plays a double-role (that is, that the same experience token plays both an enabling and an evidential role) in alleged instances of a priori knowledge of truths expressed by colour inclusion sentences would still allow defenders of the apriority of this kind of truth to resort to a modified conception of a priori truth according to which, even though experience never plays a merely enabling role in instances of knowledge of the kind of truth in question, the latter should still be regarded as a kind of a priori truth. That is, one might suspect that some plausible conception of a priori truth appealing, in some way, to the notion of a double-role can be used to vindicate the apriority of the kind of truth in question, without resulting in an inflationary use of the label ‘a priori truth’. Unfortunately for defenders of the apriority of the kind of truth in question, however, no such conception of a priori truth is available. The most self-suggesting candidate for the desired kind of modification of the traditional conception of a priori truth is one according to which it is not necessary for a truth to be a priori that it could be known without experience playing an evidential role; rather, so one might claim, it is already sufficient for a truth to be a priori that there could be instances of knowledge of it in

which no experience-token plays a *merely* evidential role. That is, one might claim that it is already sufficient for a truth to be a priori that there could be instances of knowledge of it in which all experience-tokens either play a merely enabling role or a double-role. However, as demonstrated towards the beginning of section 4.2 (see also case 4 in appendix B), knowledge of a truth that even the most passionate defenders of the a priori would be unwilling to categorize as a priori (namely, knowledge that the colour of the home jerseys of the 1986 football team of Washington State University is crimson) can be acquired without any experience-token playing a merely evidential role. Thus, the candidate in question for the desired kind of modification of the traditional conception of a priori truth overgeneralizes. This overgeneralization can be avoided by claiming that in order for a truth to be a priori, it is necessary that it could not be known without all experience-tokens that play an evidential role also playing an enabling role. However, endorsing this condition would itself overgeneralize, since not only the clear example of an a posteriori truth just mentioned, but also, as it is commonly assumed, some supposed paradigm cases of a priori truth could be known with some experience-token playing a merely evidential role. (One can, for instance, come to know an arithmetic equation by reading off the result from the display of a calculator, without the experience involved in reading off the result playing an enabling role.) I take it, therefore, that it is not at all clear how a modified conception of a priori truth, involving the notion of a double-role, could provide for a plausible vindication of the apriority of truths expressed by colour inclusion sentences. At least, the burden of proof lies with those claiming that a modification of the desired kind is available.

Third, as already mentioned in the introduction, I intend the term ‘colour inclusion sentence’ to denote all non-logically-true sentences expressing a relation of inclusion, or exclusion, of colour properties. That is, the notion of a colour inclusion sentence is supposed to cover sentences of the logical form ‘ $\forall x(Fx \rightarrow Gx)$ ’ or ‘ $\sim\exists x(Fx \& Gx)$ ’, ‘F’ and ‘G’ being colour predicate letters. It may be worth noting that my considerations against the apriority of truths expressed by sentences of this form do not constitute a case against the apriority of truths expressed by weaker modal versions of the form ‘ $\diamond\forall x(Fx \rightarrow Gx)$ ’ (e.g., ‘It is possible that all crimson things are red (all over)’) and ‘ $\diamond\sim\exists x(Fx \& Gx)$ ’ (e.g., ‘It is possible that nothing is both red (all over) and blue (all over)'), since the truth of what is expressed by such sentences already follows from the fact that there are possible worlds (or possible states of the world) not containing any crimson, red, or blue things. And my considerations in the foregoing chapters certainly

do not constitute a case against the apriority of this latter fact. However, my considerations in the foregoing chapters *do* constitute a case against the apriority of truths expressed by sentences of the form ‘ $\diamond\exists x(Fx \ \& \ Gx)$ ’ (e.g., ‘It is possible for something crimson (all over) to be red (all over)’) or $\diamond\exists x(Fx \ \& \ \sim Gx)$ (e.g., ‘It is possible for something red (all over) not to be blue (all over)’ – truths that might be regarded as even better candidates for counting as a priori than those expressed by sentences like ‘All crimson things are red’ and ‘Nothing is both red (all over) and blue (all over)’.

Fourth, and finally, I would like to present an example of how the general kind of argumentation presented in the foregoing chapters might be of some use outside the realm of colour inclusion. Mišćević 2005a argues that issues about the apriority or aposteriority of propositions expressed by sentences like ‘Whales are animals’ are context-sensitive, such that the propositions in question are a priori from the perspective of certain contexts of ascription, and a posteriori from the perspective of certain other such contexts. He further claims that the propositions in question count as a posteriori according to standards of “a more demanding context of ascription” (p. 72), and that the propositions in question should thereby be regarded as “deeply *a posteriori* and superficially *a priori*” (p. 73).⁴² His reason for this claim concerns the kind of justification that people normally provide (if asked to) for asserting that whales are animals:

“Offering analytic justification for our belief is not how we normally proceed with the question whether whales are animals. We normally point out that all the observed whales share all the important traits of animals, and that there are strong reasons to believe that all whales do. It would not be inappropriate to take one’s interlocutor to the Museum of natural history, and show her or him the actual specimens, and also some documentaries on whales.” (2005a, p. 70; a very similar formulation also appears in his 2005b, p. 305.)

Since Mišćević takes the supposed fact that it would not be inappropriate to justify the truth in question in the way described to reveal the deep aposteriority of this truth, it is to be expected that he regards the experience involved in observing actual specimens of whales and watching documentaries on them as playing an evidential role in the

⁴² Despite his claim that propositions expressed by sentences like ‘Whales are animals’ are deeply a posteriori, Mišćević nevertheless regards them as analytic, thereby challenging the “commonplace in epistemology that any analytic proposition [...] is a priori” (2005a, p. 55). Apart from the fact that analyticity is, for good reasons, commonly taken to be a property of sentences rather than a property of propositions (see, for instance, Russell 2008, pp. 21f.), I regard his argumentation against the supposed commonplace just mentioned – an argumentation presented again in his 2005b and 2006 – as unconvincing. However, I will not deal with this issue here, since it does not concern the point I wish to make.

interlocutor's acquisition of the piece of knowledge in question. Accordingly, he seems to assume that the experience in question does not merely provide for the interlocutor's understanding or grasp of the truth in question. However, he does not give any reason for supposing that taking someone who is curious or uncertain about whether whales are animals to a museum and showing her documentaries should not simply count as a lavish way of teaching her the meaning of the term 'whale'. Mišćević's case for the deep aposteriority of the truth expressed by 'Whales are animals' should therefore be regarded as incomplete.

I do not claim that an argument of the kind exemplified by the argument from analogy can provide a convincing case for the aposteriority of the truth expressed by 'Whales are animals'. Justifying such a claim would require a lot of additional argumentation, analogous to my argumentation in favour of premise 1 of the argument from analogy. Rather, what I would like to suggest is that an argument of the kind exemplified by the argument from analogy – or at least the first half of such an argument – might provide the missing explanation of why taking someone who is curious or uncertain about whether whales are animals to a museum and showing her documentaries could count as not simply teaching her the meaning of one of the relevant terms. This explanation would be along the following lines. Analogously to Norman in the scenario construed in section 1.4, one might come to understand the terms 'animal' and 'whale' from at least two different angles. That is, one might come to understand the term 'animal' by learning about the (basic) biological differences between animals and, for instance, plants. And one might, so the argument would continue, come to understand the term 'whale' by learning to distinguish whales from non-whales on the basis of certain superficial criteria, without knowing whether whales share the biological features one associates with the term 'animal'. Thus, one would conclude that one could understand the relevant terms (and have good offline reasoning skills corresponding to one's way of understanding the terms) without being in a position to know that whales are animals. By visiting a museum of natural history and watching documentaries about whales, one can learn that the things one is able to identify as whales share the biological features one associates with the term 'animal', and one thereby comes to know that whales are animals.

This way of appealing to an argument of the kind exemplified by the argument from analogy would, if the argument in question is sound, provide the missing explanation of why the experience involved in the scenario described by Mišćević could

count as playing an evidential, rather than a merely understanding-providing role in the interlocutor's acquisition of the piece of knowledge in question. Moreover, if no sound argument of the kind just sketched were available, it would not be clear how to provide *any* explanation of why the experience involved in the scenario described by Mišćević could count as playing an evidential role. Thus, even though, as already mentioned, I do not claim that the kind of argument just sketched is sound, I do claim that the soundness of this kind of argument is what Mišćević's case for the deep aposteriority of the truth expressed by 'Whales are animals' depends on. Thus, considering an argument analogous to (the first half of) the argument from analogy would be useful in evaluating the prospects of Mišćević's argumentation.

It may well turn out that considering an argument of the kind sketched above will not provide a clear answer to questions about the apriority or aposteriority of truths expressed by sentences like 'Whales are animals'. That is, it may turn out to be a vague matter whether someone able to distinguish whales from non-whales on the basis of certain superficial criteria, without knowing (or believing) that whales are animals, can count as understanding the term 'whale'.⁴³ (Compare Hilary Putnam's uncertainty regarding whether, if the things called 'cats' had always been robots, one could express a truth by uttering 'Cats turned out not to be animals' without changing the meaning of some of the relevant terms (see his 1962, pp. 660f.)) However, I take it that such a result would not speak against the usefulness of applying the kind of argument in question. Rather, it would seem reasonable to take such a result as a useful indicator that the apriority or aposteriority of truths expressed by sentences like 'Whales are animals' is itself a vague matter.

⁴³ Matters are clearer regarding 'Whales are mammals', as plausibly assumed by Donnellan 1962: "A sailor may have the ability to spot whales with great accuracy; he may know their habits and migrations; [...]" (p. 649), and he adds: "[...] I think, we would say that our sailor knew what the word 'whale' meant even though he was ignorant of the fact that whales are mammals." (p. 651)

APPENDIX A

The argument from analogy (referring to the modified version of the scenario involving Norman, described on pp. 23f.):

- (P1) In virtue of possessing the reliable skill of correctly applying the word ‘crimson’ on the basis of observing the surfaces of the objects in question through a microscope, Norman understands the word ‘crimson’.
- (P2) If some experiences necessary for a certain way of coming to know a certain truth play neither the role of providing for the subject’s understanding of the respective sentence nor the role of providing for the subject’s offline reasoning skills corresponding to the subject’s way of understanding the respective sentence, then these experiences play an evidential role in the subject’s knowledge of that truth.

From P1 and P2, together with the description of the modified version of the scenario involving Norman, the following intermediate conclusion can be inferred:

- (C1) In the modified version of the scenario, some of Norman’s experiences regarding the superficial looks of crimson things play an evidential role in his knowledge of R1.

The argument continues as follows:

- (P3) If C1 holds, then Norman’s experiences regarding the superficial looks of crimson things also play an evidential role in his knowledge of R1 in the original version of the scenario.

From P3 and C1, the following can be inferred:

- (C2) In the original version of the scenario, Norman’s experiences regarding the superficial looks of crimson things play an evidential role in his knowledge of R1.

APPENDIX B

Case 1 (introduced on pp. 88f., as a counterexample to P3_{g1}):

	Susan	Anna
how she comes to understand the term 'crimson'	During a school lesson, she looks at various colour samples and is told by the teacher which of them count as crimson.	She looks at the photos attached to the newspaper report (without reading either the article or the caption), and is told that the colour of the jerseys is one of the paradigm samples of crimson.
how she comes to know that the home jerseys of the 1986 football team of Washington State University were crimson	She looks at the photos attached to the newspaper report and reads the caption.	She relies on oral testimony.

Case 2 (introduced on p. 91, as a counterexample to P3_{g2}):

	Susan	Anna
how she comes to understand the term 'crimson'	She is told that the colour of the home jerseys of the 1986 football team of Washington State University is crimson, and then looks at an original exemplar of one of the jerseys in a football museum.	same as in case 1
how she comes to know that the home jerseys of the 1986 football team of Washington State University were crimson	After having forgotten about the fact that the colour of the home jerseys of the 1986 football team of Washington State University is crimson, she looks at the photos attached to the newspaper report and reads the caption.	same as in case 1

Case 3 (introduced on pp. 95f., as a counterexample to P3_{g4}):

	Susan	Anna
how she comes to understand the term 'crimson'	same as in case 1	same as in cases 1 and 2
how she comes to know that the home jerseys of the 1986 football team of Washington State University were crimson	She relies both on the experience described in case 1 and on oral testimony.	same as in cases 1 and 2

Case 4 (introduced on p. 102, as a plausible candidate for an example which involves an experience-token that plays both an enabling and an evidential role in the same instance of knowledge; also appealed to in the conclusions as an example of an instance of a posteriori knowledge in which all experience-tokens that play an evidential role also play an enabling role):

	Susan	Anna
how she comes to understand the term 'crimson'	irrelevant	She looks at the photos attached to the newspaper report, reads the caption, and is told that the home jerseys of the 1986 football team of Washington State University were crimson.
how she comes to know that the home jerseys of the 1986 football team of Washington State University were crimson	irrelevant	She comes to acquire the relevant piece of knowledge in the course of coming to understand the term 'crimson'.

REFERENCES

- Achinstein, P. (1968), *Concepts of Science. A Philosophical Analysis*. Johns Hopkins Press
- Audi, R. (1998), *Epistemology. A Contemporary Introduction to the Theory of Knowledge*. Routledge
- Bealer, G. (2003), "Propositions". *Mind* 107, pp. 1-32
- Block, N. (1999), "Sexism, Racism, Ageism, and the Nature of Consciousness". *Philosophical Topics* 26, pp. 39-70
- Boghossian, P. A. (2011), "Williamson on the A Priori and the Analytic". *Philosophy and Phenomenological Research* 82, pp. 488-497
- BonJour, L. (1998), *In Defense of Pure Reason. A Rationalist Account of A Priori Justification*. Cambridge University Press
- BonJour, L. (2005), "In Defense of the A Priori". In: Steup, M. & Sosa, E. (eds.), *Contemporary Debates in Epistemology*. Blackwell, pp. 98-105
- Byrne, D. & Hilbert, D. R. (2003), "Authors' Response". *Behavioral and Brain Sciences* 26, pp. 52-59
- Campbell, J. (1993), "A Simple View of Colour". In: Haldane, J. & Wright, C. (eds.), *Reality, Representation, and Projection*. Oxford University Press, pp. 257-268
- Cassam, Q. (2000), "Rationalism, Empiricism, and the A Priori". In: Boghossian, P. & Peacocke, C. (eds.), *New Essays on the A Priori*. Oxford University Press, pp. 43-64
- Casullo, A. (2003), *A Priori Justification*. Oxford University Press
- Chalmers, D. J. (1996), *The Conscious Mind. In Search of a Fundamental Theory*. Oxford University Press
- Chalmers, D. J. (2002), "On Sense and Intension". *Philosophical Perspectives* 16, pp. 135-182
- Chalmers, D. J. (2004), "Epistemic Two-Dimensional Semantics". *Philosophical Studies* 118, pp. 153-226
- Chalmers, D. J. (2006), "The Foundations of Two-Dimensional Semantics". In: García-Carpintero, M. & Macià, J. (eds.), *Two-Dimensional Semantics*. Oxford University Press, pp. 55-140

- Chalmers, D. J. & Jackson, F. (2001), "Conceptual Analysis and Reductive Explanation". *The Philosophical Review* 110, pp. 315-360
- Churchland, P. (2007), "On the Reality (and Diversity) of Objective Colors: How Color-Qualia Space Is a Map of Reflectance-Profile Space". *Philosophy of Science* 74, pp. 119-149
- Cohen, J. (2004), "Color Properties and Color Ascriptions: A Relationalist Manifesto". *The Philosophical Review* 113, pp. 451-506
- Cohen, J. (2009), *The Red and the Real. An Essay on Color Ontology*. Oxford University Press
- Cohen, J. (2012), "Redder and Realer: Responses to Egan and Tye". *Analytic Philosophy* 53, pp. 313-326
- Cohen, J., Hardin, C. L. & McLaughlin, B. P. (2006), "True Colours". *Analysis* 66, pp. 335-340
- Cohen, J., Hardin, C. L. & McLaughlin, B. P. (2007), "The Truth about 'The truth about true blue'". *Analysis* 67, pp. 162-166
- Cornsweet, T. N. (1970), *Visual Perception*. Academic Press
- Crimmins, M. (1992), "Context in the Attitudes". *Linguistics and Philosophy* 15, pp. 185-198
- Crimmins, M. & Perry, J. (1989), "The Prince and the Phone Booth: Reporting Puzzling Beliefs". *The Journal of Philosophy* 86, pp. 685-711
- Davies, M. & Humberstone, L. (1980), "Two Notions of Necessity". *Philosophical Studies* 38, pp. 1-30
- Donnellan, K. S. (1962), "Necessity and Criteria". *The Journal of Philosophy* 59, pp. 647-658
- Evans, G. (1979), "Reference and Contingency". *The Monist* 62, pp. 161-189
- Field, H. (2000), "Apriority as an Evaluative Notion". In: Boghossian, P. & Peacocke, C. (eds.), *New Essays on the A Priori*. Oxford University Press, pp. 117-149
- Fitch, G. W. (1984), "Two Aspects of Belief". *Philosophy and Phenomenological Research* 45, pp. 87-101
- Frege, G. (1948), "Sense and Reference". *The Philosophical Review* 57, pp. 209-230. Transl. from "Über Sinn und Bedeutung". *Zeitschrift für Philosophie und philosophische Kritik* 100, 1892, pp. 25-50
- Gert, J. (2008), "What Colors Could Not Be: An Argument for Color Primitivism". *The Journal of Philosophy* 105, pp. 128-155

- Glüer, K. (2007), "Colors without Circles?". *Erkenntnis* 66, pp. 107-131
- Grimm, S. R. (2006), "Is Understanding a Species of Knowledge?". *British Journal for the Philosophy of Science* 57, pp. 515-535
- Gross, S. (2005), "Linguistic Understanding and Belief". *Mind* 114, pp. 61-66
- Hale, B. (1987), *Abstract Objects*. Blackwell
- Hardin, C. L. (1988), *Color for Philosophers. Unweaving the Rainbow*. Hackett
- Jackson, F. (1982), "Epiphenomenal Qualia". *The Philosophical Quarterly* 32, pp. 127-136
- Jackson, F. (1986), "What Mary Didn't Know". *The Journal of Philosophy* 83, pp. 291-295
- Jackson, F. (1996), "The Primary Quality View of Color". *Philosophical Perspectives* 10, pp. 199-219
- Jackson, F. (1998), *From Metaphysics to Ethics. A Defence of Conceptual Analysis*. Clarendon Press
- Jackson, F. (2004), "Why We Need A-Intensions", *Philosophical Studies* 118, pp. 257-277
- Jenkins, C. S. (2008), "Modal Knowledge, Counterfactual Knowledge and the Role of Experience". *The Philosophical Quarterly* 58, pp. 693-701
- Jenkins, C. S. I. (2013), "Naturalistic Challenges to the A Priori". In: Casullo, A. & Thurow, J. C. (eds.), *The A Priori in Philosophy*. Oxford University Press, pp. 274-290
- Jenkins, C. S. I. & Kasaki, M. (forthcoming), "The Traditional Conception of the A Priori". *Synthese*
- Kalderon, M. E. (2007), "Color Pluralism". *Philosophical Review* 116, pp. 563-601
- Kant, I. (1998), *Critique of Pure Reason*. Transl. and ed. by Guyer, P. & Wood, A. W. Cambridge University Press
- Kaplan, D. (1989), "Demonstratives. An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals". In: Almog, J., Perry, J. & Wettstein, H. (eds.), *Themes from Kaplan*. Oxford University Press, pp. 481-563
- Kitcher, P. (1984), *The Nature of Mathematical Knowledge*. Oxford University Press

- Kripke, S. A. (1979), "A Puzzle about Belief". In: Margalit, A. (ed.), *Meaning and Use. Papers Presented at the Second Jerusalem Philosophical Encounter April 1976*. Reidel, pp. 239-283
- Kripke, S. A. (1980), *Naming and Necessity*. Harvard University Press
- Kvanvig, J. L. (2003), *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press
- Larson, R. K. & Ludlow, P. (1993), "Interpreted Logical Forms". *Synthese* 95, pp. 305-355
- Lewis, D. (1968), "Counterpart Theory and Quantified Modal Logic". *The Journal of Philosophy* 65, pp. 113-126
- Lewis, D. (1971), "Counterparts of Persons and Their Bodies". *The Journal of Philosophy* 68, pp. 203-211
- Lewis, D. (1981), "What Puzzling Pierre Does Not Believe". *Australasian Journal of Philosophy* 59, pp. 283-289
- Longworth, G. (2008), "Linguistic Understanding and Knowledge". *Noûs* 42, pp. 50-79
- MacFarlane, J. (2005), "The Assessment Sensitivity of Knowledge Attributions". *Oxford Studies in Epistemology* 1, pp. 197-233
- MacFarlane, J. (2007), "Semantic Minimalism and Nonindexical Contextualism". In Preyer, G. & Peter, G. (eds.), *Context-Sensitivity and Semantic Minimalism. New Essays on Semantics and Pragmatics*. Oxford University Press, pp. 240-250
- Malkoc, G., Kay, P. & Webster, M. A. (2005), "Variations in Normal Color Vision. IV. Binary Hues and Hue Scaling". *Journal of the Optical Society of America A* 22, pp. 2154-2168
- McLaughlin, B. P. (2003a), "Colour, Consciousness, and Colour Consciousness". In: Smith, Q., & Jokic, A. (eds.), *Consciousness. New Philosophical Perspectives*. Clarendon Press, pp. 97-154
- McLaughlin, B. P. (2003b), "The Place of Colour in Nature". In: Mausfeld, R. & Heyer, D. (eds.), *Colour Perception. Mind and the Physical World*. Oxford University Press, pp. 475-502
- Miščević, N. (2005a), "Is Apriority Context-Sensitive?". *Acta Analytica* 20, pp. 55-80
- Miščević, N. (2005b), "Empirical Concepts and A Priori Truth". *Croatian Journal of Philosophy* 5, pp. 289-315
- Miščević, N. (2006), "Devitt's Shocking Idea and Analyticity without Apriority". *Croatian Journal of Philosophy* 6, pp. 69-96

- Neitz, J. & Jacobs, G. H. (1986), "Polymorphism of the Long-Wavelength Cone in Normal Human Colour Vision". *Nature* 323, pp. 623-625
- Neitz, J., Neitz, M. & Jacobs, G. H. (1993), "More than Three Different Cone Pigments among People with Normal Color Vision". *Vision Research* 33, pp. 117-122
- Neitz, M., Neitz, J. (1998), "Molecular Genetics and the Biological Basis of Color Vision". In: Backhaus, W. G. K., Kliegl, R. & Werner J. S. (eds.), *Color Vision. Perspectives from Different Disciplines*. Walter de Gruyter, pp. 101-119
- Peacocke, C. (1983), *Sense and Content. Experience, Thought, and Their Relations*. Clarendon Press
- Peacocke, C. (1984), "Colour Concepts and Colour Experience". *Synthese* 58, pp. 365-381
- Peacocke, C. (2000), "Explaining the A Priori: The Programme of Moderate Rationalism". In: Boghossian, P. & Peacocke, C. (eds.), *New Essays on the A Priori*. Clarendon Press, pp. 255-285
- Pettit, D. (2002), "Why Knowledge Is Unnecessary for Understanding Language". *Mind* 111, pp. 520-550
- Powell, L. (2012), "How to Refrain from Answering Kripke's Puzzle". *Philosophical Studies* 161, pp. 287-308
- Putnam, H. (1956), "Reds, Greens, and Logical Analysis". *The Philosophical Review* 65, pp. 206-217
- Putnam, H. (1962), "It Ain't Necessarily So". *The Journal of Philosophy* 59, pp. 658-671
- Putnam, H. (1975), "The Meaning of 'meaning'". *Mind, Language and Reality. Philosophical Papers, vol. 2*. Cambridge University Press, pp. 215-271 (originally published in Gunderson, K. (ed.), *Language, Mind and Knowledge. Minnesota Studies in the Philosophy of Science, vol. 7*. University of Minnesota Press, 1975, pp. 131-193)
- Richard, M. (2004), "Contextualism and Relativism". *Philosophical Studies* 119, pp. 215-242
- Russell, B. (1912), *The Problems of Philosophy*. Hazell, Watson and Viney
- Russell, G. (2008), *Truth in Virtue of Meaning. A Defence of the Analytic/Synthetic Distinction*. Oxford University Press
- Salmon, N. U. (1986), *Frege's Puzzle*. MIT Press
- Schiffer, S. (1977), "Naming and Knowing". *Midwest Studies in Philosophy* 2, pp. 28-41

- Schiffer, S. (1994), "A Paradox of Meaning". *Noûs* 28, pp. 279-324
- Soames, S. (1987), "Direct Reference, Propositional Attitudes, and Semantic Content". *Philosophical Topics* 15, pp. 47-87
- Soames, S. (1999), *Understanding Truth*. Oxford University Press
- Stalnaker, R. C. (1976), "Possible Worlds". *Noûs* 10, pp. 65-75
- Stalnaker, R. C. (1978), "Assertion". In: Cole, P. (ed.), *Syntax and Semantics. Vol. 9: Pragmatics*. Academic Press, pp. 315-332
- Stalnaker, R. C. (1984), *Inquiry*. MIT Press
- Stanley, J. (1999), "Understanding, Context-Relativity, and the Description Theory". *Analysis* 59, pp. 14-18
- Stanley, J. (2005), *Knowledge and Practical Interests*. Clarendon Press
- Summerfield, D. M. (1991), "Modest A Priori Knowledge". *Philosophy and Phenomenological Research* 51, pp. 39-66
- Tye, M. (2006a), "The Puzzle of True Blue". *Analysis* 66, pp. 173-178
- Tye, M. (2006b), "The Truth about True Blue". *Analysis* 66, pp. 340-344
- Tye, M. (2012), "Cohen on Color Relationism". *Analytic Philosophy* 53, pp. 297-305
- Williamson, T. (2006), "Conceptual Truth". *Proceedings of the Aristotelian Society Supplementary Volume* 80, pp. 1-41
- Williamson, T. (2007), *The Philosophy of Philosophy*. Blackwell
- Williamson, T. (2011), "Reply to Boghossian". *Philosophy and Phenomenological Research* 82, pp. 498-506
- Williamson, T. (2013), "How Deep is the Distinction between A Priori and A Posteriori Knowledge?". In: Casullo, A. & Thurow, J. C. (eds.), *The A Priori in Philosophy*. Oxford University Press, pp. 291-312