A Priori truth in the natural world : a non-referentialist response to Benacerraf's dilemma

Novák, Z.

CHAPTER 7

The Non-Referentialist Alternative: A Representationist Construal of *A Priori* Truth

Introduction

In the previous three chapters, I argued that a proper construal of truth in the semantics of our purportedly *a priori* discourses about abstract domains must be realist but not referentialist (i.e. platonist) in character. The chief objection to the alternative, *prima facie* viable, non-realist and/or referentialist construals was that they cannot explain some relatively obvious features of our cognitive and linguistic practice in discourses about abstract domains.

A reaction to be expected from the advocates of the rejected deflationist, anti-realist or platonist positions would be to argue that the semantical framework proposed by their realist and non-referentialist opponents cannot meet the explanatory requirements set for an account of the relevant truths either. If this charge turned out to be true, then there would be no more reason for adopting the latter perspective, than any other of those examined before.[1]

In this chapter, I shall show that this charge is unsubstantiated. In particular, I shall argue that a specific naturalist version of non-referentialist realism about truth in the

---

[1] In chapter 2, I argued that an adequate theory must provide, either in itself or in conjunction with other theories, an explanation of all "observable" (i.e. commonly recognised) characteristics of its subject matter. What this methodological principle implies is that a theory cannot be regarded as superior over its alternatives unless it can account for all observable phenomena of its subject matter.

semantics of our purportedly *a priori* discourses about abstract domains satisfies all major adequacy conditions set for such a construal in chapter 2. I shall call the version in question a *representationist* account of the relevant truths, as its core tenet is that the conditions whose (thick) obtaining or absence determines the truth value of our claims about abstract domains obtain in the realm of representations within our head, rather than in the domain of the represented abstract states of affairs.

In section 1, I shall argue that, despite the wide consensus among present-day philosophers about the correctness of a general referentialist construal of truth, in view of our actual cognitive and linguistic practice, we have no reason to suppose that there is a conceptual link between our notion of truth conditions and our notion of subject matters or intended referents, which would make the idea of a non-referentialist construal of certain truths conceptually objectionable.

In section 2, I shall outline an ontologically naturalist account of how we can develop truth-apt representations about abstract domains with non-referential truth conditions. The account is meant to illuminate some details of the representationist construal advocated in this work, and demonstrate that the conception satisfies the third adequacy condition set for an account of truth in the semantics of discourses about abstract domains in chapter 2.

In section 3, I shall confront the representationist construal with the other explanatory challenges specified in chapter 2, and show that, in contrast with the referentialist accounts discussed earlier, this version of non-referentialism provides us not merely with a suitable response to Benacerraf's updated and generalised dilemma in the philosophy of discourses about abstract domains, but also with acceptable accounts of all those phenomena whose joint explanation we agreed to regard as a minimal condition of adequacy for a construal of the relevant truths.

Finally, in section 4, I shall briefly recall the original broader perspective of our investigation and suggest that the representationist conception advocated here qualifies as an adequate characterisation of the nature of *a priori* truth, and

therewith the nature of *a priori* knowledge and justification, in general.

1. The Consistency of the Non-Referentialist Alternative

Misgivings at the classical rationalist doctrines of *a priori* knowledge of abstract domains, according to which human beings can discover facts of platonic realms by the competent use of reason, resulted in various anti-realist or naturalist-reductionist reactions in the history of modern philosophy. The proponents of these reactions either queried the reality of the intended abstract subject matters and with them the realist construal of the relevant *a priori* truths, or adopted a revisionist, naturalistic construal of those domains and argued that our purportedly *a priori* claims are in fact empirical in character. On the other hand, as we saw in chapter 6, those who maintain that we can acquire knowledge about platonic entities take it for granted that this knowledge is knowledge of thickly obtaining platonic conditions. A common conviction of the advocates of these alternatives is that a construal of the truth conditions of a certain class of claims amounts to a construal of the corresponding subject matters and *vice versa*. Accordingly, we cannot believe in the reality of abstract truths without believing in the thick obtaining of the intended abstract conditions, and a platonist construal of the latter amounts to a platonist construal of the former.

Consider, however, the mathematical claim that there are exactly three primes between 70 and 80. On the one hand, it seems to be clear that the subject matter of this claim is not spatiotemporal in character. It says something about some numbers, and numbers do not exist in space and time. On the other hand, it seems also clear that the claim would be true even if no one ever believed that there are exactly three primes between 70 and 80. In other words, its truth conditions seem to obtain independently of what anyone would ever think of this particular issue. So, apparently, neither the naturalist-reductionist, nor the anti-realist positions can provide us with an intuitively

plausible semantics for this simple mathematical claim. Of course, platonists gladly approve this conclusion, since they maintain that the claim is true in virtue of the thick obtaining of the platonic condition that there are three prime numbers between 70 and 80. As we have seen in chapter 6, however, platonists about abstract truths have no suitable account of how we could know that the above condition actually obtains in the suggested platonic realm.

At this point, one may wonder whether the previous strategies exhaust the conceivable doctrinal alternatives in the semantics of discourses about abstract domains. The suggestion that I shall argue for in this section is that they do not. They do not, because the shared referentialist assumption underlying these strategies is not necessarily true. It would be so if our notion of truth conditions and our notion of subject matters or intended referents were related in a way that would guarantee the identity of the notions' intended referents. Our actual cognitive and linguistic practice, however, does not seem to support the idea of the obtaining of such a relation. Truth conditions are meant to be thinkable conditions whose obtaining is necessary and sufficient for the truth of a certain truth-apt representation. Further, they are meant to be those conditions whose obtaining (or absence) we must discover in order to determine the truth value of that representation. In contrast, subject matters are particular or general states of affairs that we can think of or speak about by entertaining a thought or uttering a sentence. The former conditions may actually coincide with the latter. But determining truth values and being occasionally detected by knowing minds is not the same thing as being thought or stated to obtain within a fictive or real domain. Accordingly, there is nothing inconsistent in the idea that the conditions whose obtaining we actually detect while acquiring a piece of knowledge may differ from the subject matter of the claim expressing this knowledge.

That the non-referentialist scenario is not merely a theoretical possibility can be illustrated by the case of our knowledge of fictive states of affairs. As has been observed in chapter 6, the subject matter of the claim that Little Red Riding Hood has a grandmother is clearly non-existent. Neither the girl

nor the old lady can appear in the actual world.[2] If the truth conditions of this claim were identical with the intended fictive state of affairs, then it would be hard to understand how these conditions could obtain and how we could detect their obtaining by our cognitive faculties in the actual world. Apparently, if we want to maintain that there is something whose actual obtaining (or absence) determines the truth value of the above claim and is reliably detected by knowing minds, then we must abandon the idea that this condition is identical with the intended fictive subject matter.

If truth conditions need not be identical with subject matters or intended referents, then the opponents of classical rationalism can adopt a non-referentialist, naturalist version of realism about truth in the semantics of discourses about abstract domains without subscribing to a revisionist, naturalistic construal of the apparently abstract subject matters of these claims. They can either deny or become agnostic about the existence of platonic entities and still provide a naturalist account of our purportedly *a priori* knowledge of the intended abstract domains. In the following two sections, I shall show that the non-referentialist account under scrutiny is not merely a consistent but also a well-motivated conception of truth, at least in the semantics of our purportedly *a priori* discourses about abstract domains.

2. The Emergence of Non-Referential Truth Conditions

Having abandoned the standard referentialist construal of truth in the case of our paradigm *a priori* claims about abstract domains, one may wonder what fixes the semantic relations of such representations to their non-referential truth conditions, on the one hand, and their purportedly abstract intended referents, on the other. The development of a suitable response to these

[2] As has been emphasised in chapter 6, contrary to the case of pure logical or mathematical beliefs, our referential intentions in discourses about fictive domains guarantee the non-existence of the intended fictive subject matters.

questions about semantic content-determination is fundamental in so far as it sets the framework for subsequent accounts of other observable characteristics of the relevant sorts of representations, including those which have been listed as major *explananda* in chapter 2.

In the standard referentialist semantical framework, where truth conditions are understood in terms of intended referents, the two questions formulated above coincide: explaining what fixes the semantic relations between our truth-apt mental and physical symbols, on the one hand, and their truth conditions, on the other, is nothing else than explaining what fixes the semantic relations between those symbols and their intended referents. Moreover, since the conditions constituting the relevant referential contents are entities that we are supposed to think of or speak about, the account being sought can invoke as an *explanans* the selective work of our conscious attention as well.

In possession of such explanatory resources, the referentialist story could run briefly as follows: first, we become acquainted with particular features of the world; second, our attention singles out from among these features the most striking or practically relevant ones; third, recollectable traces of these particulars are developed and kept in our memory; fourth, relying on our recollections and recognitional abilities, we identify some contrasts and similarities among the perceived particulars, and develop conceptual representations of properties as reoccurring types or universals; fifth, we observe various temporal continuities among the particular occurrences of these types; sixth, our attention singles out the most striking or practically relevant of these continuities; seventh, recollectable traces of these particular continuities are developed and kept in our memory; eighth, relying on our recollections and recognitional abilities, we identify some of these continuities as parts of single uninterruptedly existing wholes, and develop conceptual representations of individuals as possessors of these continuously existing features as essential characteristics; finally, in declarative contexts, we try to use the acquired concepts and their linguistic expressions in line with our actual evidence as to whether or not

their identified referential declarative use conditions actually obtain.[3]

Unfortunately, in a non-referentialist semantical framework, the previous account can at most serve as an explanation of how determinate semantic relations emerge between our mental and physical symbols and their intended referents. It can not explain the emergence of such relations between those symbols and their non-referential truth conditions. Worse, it appears that an account of the latter phenomenon cannot invoke the selective work of our conscious attention as a major factor in non-referential content-determination. If truth conditions are not understood in terms of intended referents, then they can hardly be singled out from a perceived or otherwise accessed domain by our conscious attention. So, those who reject the standard referentialist construal of truth in the semantics of our paradigm *a priori* discourses about abstract domains must develop an account of how our representations about these domains become associated with their non-referential truth conditions without assuming that we ever consciously attend to or think of these conditions in the course of content-determination. My primary aim in this section is to show how an advocate of non-referentialism can meet this explanatory challenge, and account for the emergence of determinate semantic relations between our representations about abstract domains and their arguably non-referential truth conditions.[4]

---

[3] The proper elaboration and confirmation of this account must, of course, emerge from a painstaking empirical inquiry into the nature of human cognition. The present outline is meant to be merely a highly abstract and rough characterisation of what scientists might once establish of the emergence of referential contents in the natural world.

[4] As we shall see, the account will involve an explanation of the emergence of non-referential truth conditions in the case of our analytic claims in general, including those whose intended referents are to be found in the spatiotemporal world.

General Theoretical Constraints

As a point of departure, let me briefly review the most important theoretical constraints that, in line with our earlier considerations, the envisaged account of the emergence of non-referential truth conditions must observe.

In the previous chapters I endorsed the commonly accepted principle of the compositionality of semantic content. According to this principle, the semantic content of a complex representation is determined by that of the representation's components and the mode of their composition. Applied to truth conditions, the principle declares that the truth conditions of a complex representation are determined by those of the representation's basic truth-apt components and the mode of their composition. The truth conditions of the latter entities, on the other hand, are supposed to be fixed by *some* appropriate semantic correlates of their semantically most basic (conceptual or subsentential) constituents and the mode of their composition. In the standard referentialist framework, the semantic correlates in question are, of course, the intended referents of these basic representations. In a non-referentialist semantical framework, however, where truth conditions are supposed to differ from intended referents, these correlates must also differ from the subject matters of the applied concepts or expressions.

In chapter 4, I argued that truth in general can be understood as correct declarative applicability in the context of the judgement or utterance under scrutiny. As we have seen, one major advantage of this construal is that in a compositionalist semantical framework it renders our theory of truth part of our general theory of correct (declarative) symbol-application. Adopting this theoretical framework, we can say that the semantic correlates of our concepts or subsentential expressions that contribute to the determination of the truth conditions of our truth-apt thoughts and sentences are what have been called the *declarative use conditions* of these atomic representations. They are conditions whose (thick) obtaining is individually necessary and jointly sufficient for the correct applicability of the relevant

concept or expression in a certain judgement or declarative sentence in the contexts of the latter's making or utterance. In the case of claims whose truth conditions can be construed in the standard referentialist way, the declarative use conditions of the key concepts and their linguistic expressions are to be identified with the intended referents of these representations. In the case of claims, on the other hand, whose truth conditions are better understood along the suggested non-referentialist lines, the declarative use conditions of the key concepts and their linguistic expressions must not be identified with the subject matter of these contentful entities. In general terms, we can say that the truth conditions of our truth-apt representations are determined by the correct declarative use conditions of their atomic conceptual or subsentential constituents and the mode of their composition.

Of course, this result, in itself, does not resolve the explanatory puzzle concerning the emergence of non-referential truth conditions. The observation that the declarative use conditions of our concepts and subsentential expressions together with their mode of composition determine the truth conditions of our truth-apt representations does not tell us anything about how those declarative use conditions become associated with the relevant atomic representations in those cases in which they are not referential in character. In the following section, I shall address and answer this fundamental question, but before doing so, let me briefly recall a few further characteristics that, on some explanatory considerations, we can reasonably assume hold of the relevant non-referential truth conditions.

First, as has been argued in chapters 4 and 5, a metaphysically neutral (deflationist) or anti-realist theory of semantic content cannot properly explain what makes the declarative use of our contentful mental and physical symbols in the actual world objectively correct or incorrect, or our truth-apt utterances objectively true or false. Since we supposed that in the absence of this explanation no theory of the subject can be adequate, we can assume that the non-referential declarative use conditions of our paradigm *a priori* representations admit of a

realist interpretation (i.e. they obtain, if they do, independently of anyone's actual thought or knowledge of this particular circumstance).

Second, as has been shown in chapter 6, a platonist construal of these conditions would undermine the possibility of any sensible explanation of how we can know of, or reliably detect, the obtaining or absence of these conditions, and therewith the declarative applicability of the relevant symbols. So, if we want to account for the possibility of knowledge of, or reliable belief formation about, a certain domain, then we must suppose also that the declarative use conditions of our representations about that domain obtain (or not) in the epistemically accessible spatiotemporal world.

Finally, our thoughts or sentences about abstract domains are the paradigms of those representations whose endorsement or rejection is supposed to be based on *a priori* considerations or evidence. Recalling the characterisation of apriority that we adopted in chapter 1, we can also say that our knowledge, or the justification of our judgements, of the obtaining or absence of the declarative use conditions of our representations about abstract domains is supposed to be independent of our experience (i.e. the deliverances of our perception of the external world and our introspection of our own bodily states). So, if we want to account for this specific feature of our representations about abstract domains, then we must suppose that the declarative use conditions of these symbols obtain (or not) in a very specific segment of the spatiotemporal world: it cannot be external to our body, since our knowledge of this part of the world is based on perceptual evidence; and it cannot be just any part of our own body, since our knowledge of most such parts is based on introspective evidence.[5] The only part of the natural world of

[5] One could claim that any first-personal knowledge of our own bodily states is by definition introspective in character. If we adopted this terminology, then the naturalistic construal of the declarative use conditions of our claims about abstract domains suggested here would render all such claims empirical, whose justification is based on experience. As I have emphasised in the first section of chapter 1, the account of *a priori* truth and knowledge advocated in this work is fully compatible with this kind of radical empiricism. The conflict between the two positions is arguably terminological. The question is, of course, whether the empiricist acknowledges the reality of an epistemologically significant substantive contrast between what is traditionally distinguished as the paradigms of *a priori* and the paradigms of *a posteriori* truth and knowledge. What I wish to defend here is that the contrast in question is real, and that (together with some further facts about the world) it can explain all observable characteristics of the distinguished entities (viz. the *a priori* and *a posteriori* pieces of truth, knowledge, justification and evidence).

which our first-personal knowledge is conventionally not taken to be introspective, and thus empirical, is the realm of our representations within our own head. Accordingly, what I intend to suggest here is that if we want to account for the apparent apriority of our knowledge of abstract domains, then we must suppose that the declarative use conditions of our representations about these domains are some specific relations among these symbols, which may or may not obtain in our head depending on how we developed them in the course of our cognitive engagement with our natural environment.

If the previous claims about the declarative use conditions of our mental and physical symbols with abstract intended referents are correct, then what we must explain in the following section is how these particular conditions in our head become associated as declarative use conditions with those symbols without ever being singled out by our conscious attention. The explanatory task in question will be accomplished in two major steps: first, I shall show how the suggested sorts of non-referential declarative use conditions appear and become associated with some of our representations in discourses about the spatiotemporal world; second, I shall show how these conditions become subsequently associated with our representations about abstract subject matters as well.

Analyticity: The Emergence of Non-Referential Truth Conditions in Discourses about Natural Domains

An important consequence of the suggested naturalistic account of the truth conditions of our claims about abstract domains is that the determination of these non-referential contents will presuppose the development of some representations in our head about the spatiotemporal world. The representations in question are, of course, also part of the spatiotemporal world, and our epistemic access to them is realised just as much by some fallible but reliable natural mechanisms as our epistemic access to the circumstances that they purport to represent.[6] The access is realised when we successfully recall those circumstances under which the relevant symbols can be correctly applied in the declarative representational contexts under consideration. Note that this realisation does not require actual thoughts *about* these representations. When we recall those circumstances under which our concept of tree can be correctly applied in the declarative representational context 'there is a … in front of me', the intentional object of our thought is not our representation of the relevant entities, but instead the entities themselves. What this reveals is that the realm of our own representations can be epistemically accessible to us even when we are actually not thinking of this realm. It is exactly this characteristic of that realm that makes it capable of providing declarative use conditions, for some symbols in some representational contexts, that are not referential in character.

Once we have developed some representations of some aspects of the natural world, we become able to recognise whether or not the declarative use conditions of these symbols (in particular declarative representational contexts) obtain in that world. In the case of synthetic declarative representational contexts, the process is realised by two consecutive epistemic

[6] Concerning the metaphysical commitments underlying this talk of representations, see section 6 of chapter 1.

mechanisms. First, we develop access to our representations, and recall those circumstances under which they can be correctly applied in the particular declarative representational context under scrutiny. Second, we examine whether or not the recalled circumstances actually obtain in the spatiotemporal world. Since the conditions to be checked in these cases are identical with what we, by applying these representations, intend to think of or speak about, the semantics of our synthetic claims about the natural world can be described as fully referentialist in character.

In the case of our analytic claims about spatiotemporal entities, in contrast, we can establish the obtaining of the relevant declarative use conditions by means of a single epistemic mechanism. By developing access to the constituents of these complex representations and recalling the circumstances under which they can be correctly applied in each other's declarative representational contexts, we can already know whether or not the truth conditions of the analytic claims in question actually obtain. We do not need to check what obtains in the represented part of the world, because the declarative use conditions of the relevant constituents in each other's declarative representational contexts are not referential in character.

As an illustration, consider the case of the paradigm analytic claim that bachelors are unmarried men. Note that in synthetic representational contexts the declarative use conditions of the semantically basic constituents of this claim are referential in character and they are supposed to obtain (or not) in the spatiotemporal world. As a result of our cognitive activity, however, some of these basic constituents become suitably linked to each other. Due to this natural relation, when I successfully access these representations in my head, I recall that they can be correctly used in each other's declarative representational contexts independently of what obtains (beyond the detected link) in the spatiotemporal world. Since the relation whose obtaining I detect when I establish that bachelors are unmarried men is not what this analytic claim purports to be about, we can ascertain that the declarative use (or truth) conditions of this truth-apt representation are also non-referential in character.

A more dubious, but in my view correct and philosophically interesting, illustration of the existence of non-referentially true analytic representations about the spatiotemporal world is the (*prima facie* synthetic) claim that water is H_2O. It is a relatively entrenched opinion today that this identity claim expresses a synthetic and necessary truth.[7] In contrast to this opinion, I believe that all knowable necessary truths are analytic in character, and the previous claim is not synthetic and necessarily true at the same time either. The idea that it is can be explained by reference to a conflation. Those who believe that the claim expresses a synthetic necessary truth conflate two senses in which the term 'water' can be applied. The first is the ordinary (pre-scientific) sense, which has no analytic relation to our scientific notion of H_2O. If one speaks about water in this sense, then the claim that water is H_2O expresses a synthetic, contingent and approximate truth. The second sense is the more technical (scientific) one, which is introduced by definition, by invoking our notion of H_2O. If the symbol is used in this sense, then the claim expresses an analytic and necessary truth, and serves as a further example of the existence of non-referentially true analytic representations of the spatiotemporal world. The two senses might be conflated, because their substitution with each other does not alter substantially the subject matter of the symbol.

Similarly to the case of the concepts of bachelor and unmarried man, once we acquired the scientific concepts of water, hydrogen, oxygen and the relevant molecular structure, we no longer need to observe our natural environment in order to know that water is H_2O. This is because, at some stage of our cognitive development, we adopt a certain representation (viz. the scientific concept of water) and establish a certain (semantically significant) natural relation between it and our earlier established

[7] Kripke (1972/1980) has famously argued that since the symbols 'water' and 'H_2O' are both rigid designators, once their referents turn out to be identical in the actual world, they must be identical in every possible world (where this referent exists) too. Accordingly, the claim that water is H_2O must be necessarily true if it is true at all.

representations of hydrogen, oxygen and the relevant molecular structure. Due to the obtaining of this natural relation, when we successfully access these representations in our head, we recall that they can be correctly used in each other's declarative representational contexts independently of what obtains (beyond the detected links, of course) in the spatiotemporal world. Since the relation whose obtaining we actually detect when we establish that water is H_2O is not what this analytic claim purports to be about, we can ascertain that the declarative use (or truth) conditions of this truth-apt representation are not referential in character.

The main contrast between the two illustrative examples is that in the former case the association of the relevant concepts, which guarantees the analyticity and the truth of the claim in question, has no empirical motivation at all. The semantic content of our term 'bachelor' was, as a matter of fact, never independent of that of our composite term 'unmarried man'. The only reason for which we introduced the notion of bachelor to our conceptual scheme is to possess a syntactically simple mental symbol by means of which we can more simply think of unmarried men.[8] The scientific reduction of water to the chemical compound H_2O, in contrast, was a more intricate intellectual manoeuvre, where the introduction of the new scientific notion of water had more complex pragmatic motivations and some empirical preconditions as well. Since the idea that this concept is different from its ordinary counterpart and has an analytic link to the previously acquired notion of H_2O is not widely recognised, let me say a bit more about how I think the story of this reduction should be best summarised.

[8] Had we ever used our term 'bachelor', in synthetic representational contexts, without keeping track of the obtaining or absence of the synthetic declarative use conditions of our complex expression 'unmarried man', the claim that bachelors are unmarried men would have amounted, like the claims that water (in the pre-scientific sense) is H_2O and that water (in the scientific sense) is transparent, to an empirical generalisation, whose referential truth conditions may not obtain in the actual world.

With some simplification, I believe we should distinguish three stages in the development of our representation of the stuffs we have classified under the term 'water' (or its translations) over the last thousands of years. In the long first stage, the term 'water' expressed a concept whose declarative use conditions were identical with some or another group of those easily observable conditions (such as fluidity, transparency, drinkability etc.) whose joint obtaining in our natural environment was so striking and important for our daily life. At some stage of our cognitive development, however, we gradually came to appreciate the practical value of those representations whose declarative use conditions were well-determined and explanatorily more significant. Our pre-scientific notion of water did not score high enough on this scale: first, it became never entirely determined which observable "water characteristics" were actually to be taken as constituents of the declarative use conditions of our concept of water; second, none of the conceivable groupings of these characteristics constituted an explanatorily outstanding (or sufficiently significant) entity in the natural world.

In search for a more comprehensive and simpler explanation of phenomena, at the beginning of the second stage of our conceptual development, we introduced an alternative, scientific concept (a mental symbol syntactically understood), also expressed by the term 'water', whose declarative use conditions were left unspecified and to be determined by our empirical science after further inquiry into the nature of the intended referents of our pre-scientific concept of water. The mode of introduction of this new mental symbol already forecast the future emergence of an analytic link, a semantically significant natural relation, within our system of representation about the natural world. It became clear that by determining the envisaged conditions we would connect our scientific concept of water with an alternative representation of the relevant conditions.

The "discovery" of the molecular structure of water was nothing other than the accomplishment of this scientific task, which brought us to a third stage in our conceptual development.

By determining (i.e. finding an articulate representation of) the conditions whose obtaining or absence in various segments of the spatiotemporal world is meant to govern the declarative use of our scientific notion of water in synthetic representational contexts, we determined the semantic content of this mental symbol, so that our term 'water' from now on expressed two concepts with more or less determined semantic contents.

The contrast between the expressed concepts is manifest in the systematic ambiguity we can observe in the declarative application of our term 'water' in various representational contexts. When the term is meant to express our pre-scientific notion of water, the claim that water is H_2O amounts to an empirical hypothesis, which, understood as a universal generalisation, is strictly speaking false. First, stuffs called water in the pre-scientific sense of the term are not composed purely of H_2O molecules. Second, some stuffs that are composed of H_2O molecules do not qualify as water in that sense, since they do not possess all the observable characteristics we take to be essential for them to do so. Under the same interpretation, the claim that water is a transparent and drinkable liquid is not merely true, but it may be analytic, and thus necessarily true. Whether it is depends on whether or not the predicated qualities are constitutive of the declarative use conditions of the subjects' pre-scientific concept of water.[9] When the term is meant to express our scientific notion of water, on the other hand, the claim that water is H_2O proves to be analytically true, while the claim that water is a transparent and drinkable liquid amounts to an empirical hypothesis, which, understood as a universal generalisation, is strictly speaking false.

[9] Those who have an explicit idea of exactly which observable "water-characteristics" constitute the declarative use conditions of their own pre-scientific concept of water in synthetic representational contexts are supposed to have an analytic connection in their head between the latter concept and those ideas of the relevant conditions. As a general rule, a concept is always analytically related to an alternative representation of its (referential) declarative use conditions.

What we had discovered by the end of this stage was that the obtaining of the (vaguely determined) declarative use conditions of our pre-scientific notion of water, in synthetic representational contexts, strongly correlates with that of the (relatively sharply determined) declarative use conditions of our notion of the chemical molecule H_2O. It was this empirical finding which convinced us that the characteristic of being constituted of H_2O molecules could be a convenient declarative use condition for our new scientific concept of water in synthetic representational contexts. By adopting this stipulation, we ensured that the content of this new symbol became not merely well-determined and explanatorily outstanding, but also sufficiently overlapping with that of our ordinary pre-scientific notion of water. [10]

Summing up, what the previous examples illustrate is that non-referential truth conditions emerge already in our discourses about the (epistemically accessible) natural world. The claims whose truth value is apparently determined by the obtaining or absence of such conditions are our analytic claims about the spatiotemporal realm. In addition, the suggested accounts assume that the conditions in question obtain (if they do) in the domain of representations within our head. Although the emergence of these conditions is the result of our cognitive-epistemic activity, once they are established they obtain independently of our actual thoughts or beliefs about this circumstance. In fact, we can even detect these obtainings without ever thinking of the relevant

[10] Note that the role of experience in motivating this content-determining move does not imply that the analytic truths emerging with this content are empirical in character. Questions about ways of knowledge and justification can be raised only if semantic contents have already been determined. The process of content-determination is utterly conventional, subject merely to practical considerations. The establishment of truth values, in contrast, is an activity whose success or failure (providing that success in this case means holding true what is true and holding false what is false) is an entirely objective issue, determined exclusively by what obtains in the world independently of what our actual opinions are concerning this circumstance. The conflation of the two processes may result not merely in radical empiricism in epistemology, but, as I noted in chapter 5, also in anti-realist conclusions concerning the nature of truth.

representational states of affairs. Finally, the examples also illuminate that, at least in the case of discourses about spatiotemporal domains, the declarative use conditions of our symbols may be referential in some representational contexts, while non-referential in others. With these conclusions in mind, I shall turn now to the second part of the current explanation, and show how the previous representational conditions can serve as declarative use conditions for our representations about abstract domains as well.

Abstraction: The Emergence of Non-Referential Truth Conditions in Discourses about Abstract Domains

In the previous subsection, I explained how our analytic claims about the spatiotemporal world acquire their non-referential truth conditions that obtain (if they do) in the domain of representations within our head. As we saw, a crucial feature of these conditions, which enables them to play this sort of semantic role, is that we can detect their obtaining or absence without actually thinking of them. In this section, I shall argue that, in consequence of a particular concept-forming process within our head, the very same conditions can serve as non-referential truth conditions in the case of our paradigm *a priori* claims about various abstract domains as well.

As a point of departure, let me recall a specific difficulty, briefly touched upon at the beginning of this section, that may seem to balk a proper account of the emergence of non-referential truth conditions in the semantics of discourses about causally inert domains in particular. Suppose someone asks us what makes the allegedly non-referential truth conditions of our mathematical claim *that there are three prime numbers between 70 and 80* in our head the truth conditions of this particular representation. In the case of our synthetic claims about causally effective entities, the corresponding question can be answered by invoking our referential intentions and the selecting work of our conscious attention. The declarative use conditions of our

concept of white and concept of snow, for instance, are presumably singled out by some causal mechanisms including those underlying the selecting work of our conscious attention in the course of the relevant semantic content-determinations. Together with the applied rules of composition, these mechanisms can be invoked in an explanation of what makes the referential truth conditions of our claim *that snow is white* the truth conditions of this particular representation. In the case of our analytic claims about causally effective entities, where the relevant truth conditions are supposed to be non-referential in character, the same question can be answered, as we saw, by invoking our own mental activity that associates more than one of the symbols figuring in the claim in question with the same (previously determined) referential contents. One may say, for instance, that the non-referential truth conditions of our claim *that bachelors are unmarried men* are the truth conditions of this particular claim, because the mental stipulations responsible for their obtaining concerned the content of the key constituents of this particular representation (i.e. the declarative use conditions of our concept of bachelor and our concept of unmarried man).

The apparent difficulty with providing a satisfying answer to the same question in the case of our claims about causally inert domains is that here we cannot suppose that the declarative use conditions of the semantically basic constituents of these claims are singled out from the intended causally inert domains by some causal mechanism including those underlying the selecting work of our conscious attention. This is simply because there can be no causal contact between our minds and the intended causally inert domains.[11] So, if there is an appropriate account of the emergence of non-referential truth conditions in our discourses about the latter domains, then it must presumably rely on an articulate conception of how we develop representations of causally inert entities without any interaction with the relevant

[11] In chapter 6, I argued that alternative conceptions of the existence of a non-causal contact between human minds and some platonic domains are both *ad hoc* and uninformative.

intended domains, and it must also illuminate what makes a certain natural link among some representations of the natural world in our head the truth condition of a particular claim about some causally inert entities. In what follows, I shall provide an outline of such an account.

The key element that I shall rely on in this account is our ability to create new concepts with non-spatiotemporal subject matter from earlier developed ones about some spatiotemporal entities by a certain cognitive process that I shall call abstraction. Unfortunately, the terms 'abstraction' and 'abstract' are used rather ambiguously in the philosophical literature, so before advancing the promised account, I must begin with a brief specification of the sense in which I shall use these terms in that account.[12]

In one received sense, abstractness is contrasted with concreteness and it can be characterised as the property of having no autonomous ontological status. Autonomous existence is often meant to be the privilege of concrete entities. The fact that we can think of entities that have no such existence is then thought to be the consequence of a mental operation, a separating act of human mind. It is this sense in which the particular colour of a concrete object (say, a red rose) can be regarded as an abstract characteristic. The classification reflects the natural, though hardly trivial, metaphysical assumption that the colour of a particular rose exists only as a feature of a concrete individual. If the flower ceased to exist, so would the particular colour as well. We can think of this colour merely because we can separate this feature from its concrete bearer in thought by abstraction. If we were not able to carry out this operation, we could not develop the concept in question, since we could never be acquainted with the intended subject matter independently of other characteristics of the concrete particular.

Abstraction in this sense can be held to be an important operation by which we are able to create new concepts from

[12] A brief overview of the alternative notions that I shall distinguish here has been given in the fifth section of chapter 1.

some others that we have acquired earlier. We carry out this operation by singling out a proper part of the declarative use conditions of an acquired concept and regard them as the declarative use conditions of the newly introduced one. In the case of our example, we may acquire a concept of a particular rose by developing a rich, though syntactically simple, representation of a flower in our direct natural environment. Later on, even in the absence of the intended referent, we can develop a new concept of its particular colour by separating the relevant feature of the concrete object in thought and regarding it as the declarative use condition of a new mental symbol in various synthetic representational contexts. Clearly, abstraction in this sense does not lead us out of our discourses about the natural world. The new concepts developed in this manner will not cease to represent aspects of a spatiotemporal universe.

In a second sense, abstractness is contrasted with particularity and it can be characterised as the property of having no unique location in space and time or, in other words, being a universal that may be instantiated in various spatiotemporal locations. It is this sense in which redness, a property appearing in space and time, qualifies as an abstract universal. Abstractness in the second sense is clearly different from abstractness in the first. A particular colour of a concrete individual is certainly not abstract in the second sense of the term, and a universal feature characterising a number of different individuals can hardly owe its existence to any one of these concrete particulars.

Abstraction in the second sense seems also an important operation by which we can create new concepts from some others that we have acquired before. In particular, we can create concepts by which we can think of spatiotemporal features that may appear in various spatiotemporal locations. In the case of our example, after acquiring representations of particular red and green colours along the lines specified in the previous paragraph, we can simplify our representational system and create some new concepts by which we can think of and speak about these colours in general, disregarding of their particular spatiotemporal location.

Notice that abstraction in this second sense does not lead us out of our discourses about the natural world either. Although our concept of redness resulting from this operation no longer stands for a particular red colour, it nevertheless represents a universal property that can be instantiated in space and time. When we apply these concepts for identifying a subject (what we do with the concept of water in *water is H_2O*) or predicating something about a certain subject (what we do with it in *this liquid is water*), we are still supposed to be thinking of features that can appear in the spatiotemporal world.[13]

It is important to see that, in synthetic representational contexts, the declarative use conditions of our representations of abstract entities in these first two senses of the term can be construed along the standard referentialist lines. This means that an advocate of standard referentialist realism about synthetic truths need not give up her theory in the semantics of our discourses about such abstract domains.[14]

The correctness of the previous tenet is by no means conspicuous. One may wonder, for instance, how an advocate of this referentialist view would explain the emergence of the posited referential relations in the case of our synthetic claims that involve concepts of universals actually not instantiated in

[13] As it has been observed in chapter 6, the concepts whose acquisition Hale and Wright explained by reference to our ability to carry out Fregean abstraction may have subject matters that are abstract only in this second sense of the term. Despite the authors' explicit stipulation, the way they invoke Frege's Abstraction Principle in their reasoning against their anti-platonist opponents seems to suggest that the operation they have in mind does not create concepts of strictly non-spatiotemporal entities from concepts of entities that can be instantiated, and thus known by acquaintance, in space and time. See esp. fn. 28 in chapter 6.

[14] Of course, in the case of analytic claims about these domains, such as *bachelors are unmarried men*, the considerations against a referentialist account of truth presented earlier seem still adequate. If the truth conditions of these claims (again, our analytic claims involving concepts of *in re* universals instantiated in the natural world) were referential in character, then we would have no suitable explanation of how we can learn about the necessary character of these truths, and how we are able to discover them, in possession of the relevant concepts, without any check upon what obtains in their intended referential domain.

space and time (e.g. the concept of unicorn or that of a man taller than 20 feet). The question to be answered here is how we can develop an idea of the intended uninstantiated declarative use conditions without ever getting acquainted with them in the course of our cognitive development.

The most intuitive answer to this question is to suppose that our ideas of uninstantiated universals are the combinatorial results of our previously acquired ideas of instantiated universals.[15] Our idea of a unicorn, for instance, may be thought to be constructed in our head from our ideas of a body, a head, a hind leg, a tail, a horn, a forehead, a horse, a stag, a lion, being a single and being in the middle.[16]

Some may want to challenge this answer by reference to the conceivability of simple (i.e. not constructed) uninstantiated universals. According to this objection, if in synthetic representational contexts the referentialist understanding of the declarative use conditions of our concepts of spatiotemporal universals were true, then we would have no suitable explanation of how the posited referential relations could emerge in the case of our concepts of simple universals that could be but are not instantiated in space and time.

To this challenge, a referentialist may reply that our ideas of the declarative use conditions of these concepts, in synthetic

[15] A recent defence of this response can be found in Armstrong (1989). Note that Armstrong's combinatorialism is meant to be not merely an account of the *semantic* content of our ideas of uninstantiated universals, but also a *metaphysical* account of unrealised possibilities. The extension of the combinatorialist perspective from semantics to metaphysics, however, seems necessary only if we think that the factual basis of modality is to be found in the intended domains of our modal thoughts. On a non-referentialist construal of modal truths, which I propose to adopt in this work, no such extension is required or appropriate. The truth value of our claims about unrealised possibilities will be determined by some contingent facts within our head, rather than by some modal features that obtain in the intended natural domains.

[16] According to the Merriam-Webster's Online Dictionary, a unicorn is a mythical animal generally depicted with the body and head of a horse, the hind legs of a stag, the tail of a lion, and a single horn in the middle of the forehead (http://www.m-w.com/dictionary/unicorn).

representational contexts, are merely negative. What we know of them is that they are not constructible from any conditions that are instantiated in the spatiotemporal world. In the light of this minimal knowledge, these concepts can be correctly applied in any synthetic representational contexts in which this application creates a claim whose truth requires the absence of the intended alien conditions at the intended spatiotemporal locations (as in the context 'there is no … in this room').[17] Since the intended subject matter of these concepts are, *ex hypothesi*, simple uninstantiated universals (i.e. actually non-existing thinkables), a referentialist semantics in their case can account for the above minimal knowledge even in the absence of any acquaintance on our part with these alien conditions.

Having shown that the semantic content of our synthetic representations of abstract entities in the first two senses of the term can be construed along the standard referentialist lines, we can turn now to the case of those representations whose subject matter is supposed to be abstract in a third sense of the term, a sense that excludes the spatiotemporal construal of the relevant intended entities.

It is abstractness in this third sense that characterises the subject matter of our paradigm *a priori* claims and underlies the received definition of platonism as realism about abstract entities. According to this construal, abstractness is contrasted with spatiotemporality, and it can be understood as the property of

[17] Within the referentialist framework under discussion, the declarative use conditions of our concepts of simple alien universals in synthetic representational contexts cannot differ from each other. This is simply because, contrary to our concepts of instantiated and constructed uninstantiated universals, these concepts have no positive semantic link to any particular aspects of reality that they could be about. This is why we know under what circumstances we should declaratively apply our concept of unicorn in the representational context *this creature in front of me is a* …, while we have no similar knowledge in the case of concepts of simple uninstantiated universals. Nevertheless, we can develop more of the latter concepts, since their declarative use conditions can differ in analytic representational contexts. The difference can be established by convention at the time of the introduction of these new concepts into our actual conceptual framework.

having no location whatsoever, whether unique or multiple, in space and time. Beyond the intended subject matters of our paradigm *a priori* discourses, *ante rem* universals or categorical norms and values may also be classified as abstract in this sense of the term. Clearly, this third understanding is different from the previous two. Particular and universal colours appearing in space and time are not abstract in the third sense of the term, while numbers and other non-spatiotemporal individuals are not abstract in the first and the second.

Returning to our main line of thought, what I wish to suggest here is that by abstraction in this third sense we can create new concepts with non-spatiotemporal subject matters from earlier acquired ones about our spatiotemporal environment. By reference to this cognitive process we can explain how we learn determinately to refer to and make truth-apt claims about strictly non-spatiotemporal entities in the absence of a suitable causal interaction with the relevant intended domains.[18] Furthermore, the emerging relations between the respective input and output notions of this process will also explain what makes the earlier invoked natural links among our representations of the natural world in our head the truth conditions of our claims about the relevant abstract (i.e. non-spatiotemporal) domains as well.

Generally speaking, by abstraction in the third sense I mean a cognitive process by which we can create concepts of strictly non-spatiotemporal individuals from concepts of properties that can be instantiated in space and time.[19] The subject matters of two output concepts of this process are meant to differ exactly when the subject matters of the corresponding input concepts

[18] The account is not meant to explain the emergence of our concepts of causally inert subject matters, like values and normative properties, that are supposed to appear in the spatiotemporal world.

[19] As noted in chapter 6, if the stipulation that the instances of Frege's Abstraction Principle provide us with implicit definitions of concepts that are about strictly non-spatiotemporal objects is taken seriously with all its consequences, then the principle can be regarded as a formally adequate characterisation of abstraction in the currently intended third sense of the term.

differ too. Further, the process is meant not to affect the analytic relations of the input concepts. Consequently, any two concepts created by this process will be analytically related if and only if their antecedents were analytically related too. The correspondence is guaranteed by the fact that the operation keeps the declarative use conditions of the output concepts in each others' analytic representational contexts identical with those of the respective input concepts in each others' analytic representational contexts. Due to this identity, the truth conditions of our analytic claims composed of the output concepts will be also identical with those of our analytic claims composed of the respective input concepts of this operation. In principle, we can conceive the abstract counterpart of any property that can be instantiated in space and time.[20]

In synthetic representational contexts, such as '… exists in the non-spatiotemporal part of the world', the declarative use conditions of the resulting concepts preserve their referential character. Accordingly, the above explanation of the emergence of our concepts of abstract entities explains the emergence of our ideas of these synthetic declarative use conditions as well. In chapter 6, I argued that we have no reason to suppose that we could ever discover whether or not these referential conditions actually obtain. Consequently, I believe that we cannot hope to acquire synthetic knowledge of non-spatiotemporal domains. Of course, the impossibility of such knowledge does not imply that we cannot acquire any sort of knowledge of these domains. This is because our ideas of abstract entities in the third sense can be declaratively applied in analytic representational contexts as well, in which the conditions whose obtaining or absence determines whether or not we can correctly apply the relevant concepts in those contexts are non-referential in character.

The only prerequisite of the existence of such representational contexts is that at least some of the actual input notions of abstraction in the third sense must be analytically

[20] Plato's claim that every spatiotemporal characteristic is merely the reflection of an atemporal Form provides the clearest illustration of this capacity.

related to each other. Since the operation does not affect these analytic relations, any concepts which have been developed by abstraction in the third sense from analytically related notions of spatiotemporal entities will be analytically related too. If they appear in each other's truth-apt representational contexts, then our reliable epistemic access to the relevant representational facts in our head will (mostly) enable us to establish the truth value of these claims. We will not need to check what obtains in the represented abstract part of the world, because the declarative use conditions of the relevant concepts in each other's representational contexts will be non-referential in character. As we observed before, although the emergence of these conditions is the result of our cognitive-epistemic activity, once they are established they obtain independently of our actual thoughts or beliefs about this circumstance, and we can detect these obtainings without ever thinking of the relevant representational states of affairs.

Consider, for instance, the acquisition of numerical concepts applied in pure arithmetic, such as the concept of number three. The subject matter of this concept is an entity that cannot appear in space and time. It cannot appear there, because we do not think of it as something possibly spatiotemporal.[21] If my

[21] To be more exact, our *referential intentions* accompanying the use of this concept exclude the viability of the naturalist construal of this subject matter. If we maintained that number three is non-spatiotemporal merely because we never saw this object in the natural world, then our thought would not guarantee the abstractness of this entity in the relevant sense of the term. One may object that if the previous claim were true, then the sentence 'The number of chairs in front of me is three, just like the number of primes between 70 and 80, and the number of King Lear's daughters in Shakespeare's drama' could not be true. It could not be true, since a strictly non-spatiotemporal entity cannot be identical with a spatiotemporal one appearing in the real world (as the number of chairs in front of me) or in a fictive universe (as the number of King Lear's daughters in Shakespeare's drama). Since the above sentence can be true, our referential intentions accompanying the use of our concept of number three must allow for a domain-independent construal of the concept's subject matter. Note, however, that if we are asked whether the arithmetic object number three could appear in space and time, our answer is unanimously negative. This supports the original claim that the way we think of the objects of pure arithmetic excludes the

considerations against platonist contact theories of knowledge of causally inert entities advanced in chapter 6 are sound, then we must abandon the idea that the semantic content of this concept is determined by the interaction of our conscious attention with the domain of pure arithmetic. On the other hand, we clearly possess the concept and we use it in various representational contexts in a highly disciplined manner. Apparently, this use is informed, under normal epistemic circumstances, by the obtaining (or absence) of some non-epistemic conditions in the actual world that we can reliably detect by a purportedly *a priori* epistemic mechanism. So, how did we acquire this concept? How did this notion acquire its strictly non-spatiotemporal subject matter and its arguably spatiotemporal declarative use conditions in those representational contexts in which we are supposed to apply it in an evidence-governed way?

According to the account being proposed here, our concept of number three is developed by abstraction in the third sense from our concept of the number of objects in triples or, what is

naturalist construal of these subject matters. The apparent conflict between the quoted examples, however, disappears if we recall that in natural language a physical symbol can be applied to express various mental representations. The fact that, under some actual circumstances, our numerical expression 'three' can be correctly applied in the representational context '*The number of chairs in front of me is* …' would be in conflict with the fact that the same term can be correctly applied in the representational context '… *cannot appear in space and time*' as well only if it were supposed that the term expresses the same concepts in these applications. (For a currently irrelevant complication, see also fn. 27 in chapter 6.) Dropping this assumption, we can eliminate the conflict by maintaining that the term can be correctly applied in the former context, because it expresses there a concept of a numerical property that may characterise groups of individuals independently of whether these are fictive or real, or abstract or spatiotemporal, and it can be correctly applied in the latter context, because it expresses there a concept of an individual that cannot appear in space and time. What I wish to explain in the main text is the emergence of the semantic content of the latter notion and *its* linguistic expression. The conflation of the two senses in which mathematical symbols can be applied (viz. the one observed in pure mathematics and the other operative in the empirical sciences) is at the heart of the influential empiricist reasoning from the indispensability of mathematics in the empirical sciences to the existence of abstract mathematical objects and properties.

the same, from our concept of being a triple (i.e. a numerical property that characterises spatiotemporal groups with three members).[22] The input concept of this operation is supposed to be acquired, together with the concepts of some other basic numerical properties and operations, via acquaintance with the spatiotemporal instances of the intended entities.[23]

The process fully determines the subject matter of our notion of number three. It makes us apply this notion with the intention to think of an individual that cannot appear in space and time and is the strictly non-spatiotemporal correlate of the numerical property of being a triple. This much specificity in our referential intention accompanying the declarative use of our concept of number three is sufficient to distinguish this subject matter from the intended referents of any other notions developed by this type of operation, and, consequently, from any other entity that we can think of.

How about the determination of the respective declarative use conditions? Well, in synthetic representational contexts, the declarative use conditions of our concept of number three are supposed to preserve their referential character. Accordingly, by determining the previous subject matter, abstraction in the third sense can also fix the synthetic declarative use conditions of our

[22] Here I suppose that the semantic content of the alternative input expressions invoked is the same, so they express the same mental symbol or the same concept. By applying them declaratively in various representational contexts, we intend to speak about a numerical property that can be multiply instantiated in space and time. The expressions in question are still not interchangeable, because they have different syntactic roles: while the term 'the number of objects in triples' can identify a certain numerical property as the subject of a proposition, the term 'is a triple' predicates this property of some subjects identified otherwise. (Again, for a minor complication concerning the intended referent of the former term, see also fn. 27 in chapter 6.)

[23] Our concepts of non-instantiated or actually unobserved numerical properties are presumably developed by recursive analytic stipulations from the most basic ones, rather than by direct acquaintance with the spatiotemporal instances of these universals. The determination of exactly which numerical concepts are acquired by acquaintance, and thus independently of each other, and which are developed from these by analytic stipulations, is a task of our empirical sciences.

concept of number three. It is the obtaining or absence of these referential conditions that we are supposed to detect in order to know, for instance, whether number three, a platonic object, exists in the actual world. Again, if my arguments against the conceivable platonist epistemologies in chapter 6 are correct, then we have no reason to suppose that we could ever discover whether or not these referential conditions actually obtain, and thus whether the synthetic claim just mentioned about number three is true or false.

Of course, in pure arithmetic our concept of number three is used in a highly disciplined and evidence-governed way. When we are asked about the truth value of an arithmetic thought involving the notion of number three, our response is supposed to be non-arbitrary, based on the recognition of some real fact, the obtaining or absence of some non-epistemic conditions in the actual world.[24] In other terms, our cognitive and linguistic practice in pure arithmetic suggests that we can still acquire some knowledge of the strictly non-spatiotemporal domain of this discipline, including the subject matter of our concept of number three.

If this knowledge cannot be synthetic, then it must be analytic in character. As we saw, the declarative use conditions of our mental and physical symbols in analytic representational contexts cannot be adequately construed in referential terms, because on such a construal we could have no suitable explanation of how we can learn about the necessary character of analytic truths, and how we can discover them, in possession of the relevant concepts, without any check of (i.e. any cognitive interaction with) what obtains in the relevant intended referential domains. By adopting a non-referentialist, naturalistic construal of the truth conditions of our standard arithmetical claims, and thus our claims about number three, we can remove the main

[24] In chapter 4, I argued that without this substantive realist construal of the declarative use conditions of our mental and physical representations, we cannot suitably explain the objectivity of their correct declarative applicability or truth.

obstacle from the path of a proper contact theory of this type of knowledge as well.

Note, however, that this move implies also that the former account of how abstraction in the third sense determines the strictly non-spatiotemporal subject matter of our concept of number three does not explain how the notion acquires its non-referential declarative use conditions in analytic representational contexts, as those occurring in pure arithmetic. Apparently, a proper account of the latter phenomenon must invoke some fact about the suggested *explanans* whose explanatory role is not limited to the determination of the relevant strictly non-spatiotemporal subject matters.

The fact that I propose to invoke at this point is, again, that abstraction in the third sense fully preserves the analytic relations of the concepts acted upon. More specifically, I claim that the operation keeps the declarative use conditions of any two output concepts in each others' analytic representational contexts identical with those of the respective input concepts in each others' analytic representational contexts. In other terms, I suggest that the non-referential truth conditions of our analytic claims involving the former concepts are the same representational conditions that constitute the non-referential truth conditions of our analytic claims involving the latter.

In pure arithmetic, for instance, we think we know that one plus one plus one equals three. In other terms, we think we have good reason to assert that our concept of number three can be correctly applied in the representational context 'One plus one plus one equals …'. The suggestion we made earlier is that we can explain the possibility of this knowledge only if we suppose that the arithmetical thought in question is analytic, and the declarative use conditions of our notion of number three in the above context are non-referential in character, obtaining in our head, in the domain of our representations of the spatiotemporal world. What needs to be explained now is what determines the alleged semantic relations between the arithmetical concept under scrutiny and these non-referential conditions. Our answer to this question is briefly that the relations are established by those

cognitive mechanisms that underlie the acquisition of our arithmetical concepts, including our idea of number three.

In the case of the above example, the explanation runs as follows. In the course of our cognitive development, first we develop our concepts of numerical properties that are instantiated in space and time. These include our concepts of being a single and being a triple. During the same time, we acquire our concept of addition and our concept of being equal, whose subject matter (an operation and a relational property, respectively) can also occur in the natural world. The emergence of the existing semantic links between these basic concepts, on the one hand, and their referential use conditions in synthetic representational contexts, on the other, can be explained by invoking our acquaintance with the spatiotemporal instances of these subject matters.

At some point in this development, maybe after the recognition of the correlation between the obtaining of the referential declarative use conditions of our concept of being a triple, on the one hand, and that of our concept of being a single and another single and yet another single, on the other, we introduce a natural link between these representations in our head, which constitutes the obtaining non-referential truth condition of various analytic thoughts involving these concepts, and becomes constitutive of the semantic content of the related notions as well.[25] By detecting the obtaining of this condition through a reliable epistemic mechanism in our head, for instance, we can know that George's single apple, Peter's single apple and John's single apple on the table in front of us constitute a triple, without checking separately, after establishing the existence of the singletons on the table, whether the referential declarative use conditions of our concept of being a triple (in the above

[25] Before the introduction of this natural link, if there is such a stage, our belief that a group of three distinct entities is a group of a single entity, another single entity, and yet another single entity amounts to an empirical generalisation, whose truth conditions are referential in character.

representational context) actually obtain in the relevant part of the spatiotemporal world.[26]

In possession of the previous concepts we can form various synthetically or analytically true ideas of the two numerical properties, the relation and the operation in question, but we cannot develop thoughts of the strictly non-spatiotemporal domain of pure arithmetic. In particular, we cannot entertain the idea that one plus one plus one equals three in the sense we are supposed to do that in pure arithmetic. In order to formulate this thought, we must acquire the concepts of number one and number three as well.[27] According to the account under consideration, these concepts are developed by abstraction in the third sense from our earlier acquired concepts of being a single and being a triple, respectively. The operation determines the subject matter, and therewith the referential declarative use conditions (for synthetic representational contexts), of the two output concepts in the manner specified above. On the other hand, it also preserves the analytic relations of the concepts it acts upon by associating its output concepts with the non-referential declarative use conditions (for analytic representational contexts) of their respective antecedents.

Due to the latter feature, our concept of number three is analytically related to our concept of one plus one plus one just like our concept of being a triple is to our concept of being a single and another singe and yet another single. Moreover, the feature also ensures that the non-referential declarative use conditions of our concept of number three in the analytic

[26] The fact that the truth of this thought requires the obtaining of some conditions in the relevant intended domain (viz. the existence of the apples denoted by the definite descriptions on the table in front of us) indicates that, despite the analytic content mentioned in the main text, the whole idea cannot be analytically true. For a thoughtful discussion of the relation of our "arithmetic" and "perceptual" criteria for the applicability of our arithmetic terms, see Craig (1975).

[27] I suppose that the concepts of addition and equality that we apply in pure arithmetic do not differ from our ideas of addition and equality in the spatiotemporal world.

representational context 'One plus one plus one equals …' are identical with the representational conditions constituting the declarative use conditions of our concept of being a triple in the representational context 'One apple plus another apple plus yet another apple make … apples on the table in front of me'. As before, the analytic links in question are constitutive of the semantic contents of the concepts involved. Accordingly, the fact that we possess the concepts of number one, number three and being equal guarantees the truth of the thought that one plus one plus one equals three, independently of whether or not the intended platonic (i.e. real and strictly non-spatiotemporal) referents of this thought actually obtain.

The account clearly removes the main obstacle from the path of a naturalist contact theory of our (purportedly *a priori*) arithmetical knowledge of number three as well. If the truth conditions of our standard arithmetical beliefs about this abstract individual obtain in the domain of representations within our heads, then we can suppose that the formation of these beliefs is constrained by those reliable epistemic mechanisms that provide us with evidence of the obtaining or absence of these conditions.

The emergence of non-referential truth conditions can be explained in similar terms in the case of our purportedly *a priori* beliefs about other strictly non-spatiotemporal domains, such as the domain of pure geometry, set theory and logic, as well. Our concepts of geometrical, set theoretic and logical objects are supposed to be developed also by abstraction in the third sense from earlier acquired concepts of properties that can be instantiated in space and time.[28] In each case, the operation

[28] Our concepts of geometrical objects, for instance, are supposed to be rooted in our concepts of spatial properties, our notions of set-theoretic objects in our notions of group-member relations, while our ideas of logical objects (e.g. platonic ideas and propositions) in our ideas of semantic properties. Note also that the account advocated here does not imply that to each concept of a strictly non-spatiotemporal object we can find in our head an earlier acquired concept of a certain property from which the former notion is developed by abstraction. Once we acquired some basic ideas of an abstract domain, we can develop further

ensures that the declarative use conditions of the relevant output concepts in analytic representational contexts coincide with those representational conditions that constitute the non-referential declarative use conditions of the corresponding input concepts in such representational contexts.

Summing up, my primary aim in this section was to show how the advocates of a non-referentialist construal of truth in the semantics of our purportedly *a priori* discourses about abstract domains can account for the emergence of determinate semantic relations between our truth-apt representations within these discourses, on the one hand, and their arguably non-referential truth conditions, on the other, without assuming that we ever consciously attend to or think of these conditions in the course of the relevant content-determination. As a first step, I provided a brief review of those theoretical assumptions that we could make of the nature of the envisaged non-referential conditions in view of some explanatory *desiderata* put forward in chapter 2. The upshot of this survey was that the conditions in question are best understood as substantive real states of affairs that obtain (if they do) in our head, in the domain of our representations that we develop in the course of our cognitive engagement with the spatiotemporal world. Second, I explained how these representational conditions may become associated, as non-referential truth conditions, with our analytic claims about the natural world. The crucial assumption that I made in this part was that the suggested representational conditions can fulfil this particular type of semantic role, because we can reliably detect their obtaining or absence without actually thinking of them. Finally, I turned to the case of our purportedly *a priori* discourses about abstract domains, and argued that, due to a particular concept-forming mechanism in our head, which generates concepts of strictly non-spatiotemporal objects from earlier acquired notions of some properties that may be instantiated in

notions of it by definition (i.e. by composing new concepts from those basic ones) as well.

space and time, the non-referential truth conditions of our analytic claims about the spatiotemporal world can serve as the non-referential truth conditions of our standard (analytic) claims about the relevant abstract domains as well.

If the account presented in this section is correct, then we can conclude that the non-referentialist conception advocated in this work can satisfy the third adequacy condition set for a theory of truth in the semantics of our paradigm *a priori* discourses about abstract domains in chapter 2. Of course, this explanatory virtue in itself does not guarantee the superiority of this conception over its referentialist alternatives. The semantical positions criticised in chapters 4, 5 and 6 are clearly not devoid of explanatory potential either. Their inadequacy became manifest only against the background of a larger pool of *explananda*. What we have learned in the previous chapters is that none of these alternatives can account, either by itself or as part of a larger theory, for all those phenomena whose joint explanation we agreed to regard as a minimal condition of adequacy for a theory of the relevant truths. Now, if we want to make a convincing case in support of the suggested non-referentialist construal of truth in the semantics of our paradigm *a priori* discourses (i.e. our discourses about abstract domains), then we must show that the conception under scrutiny can meet all explanatory requirements that we took as a minimal condition of adequacy for a theory of *a priori* truth in chapter 2. It is this task that I intend to accomplish in the section that follows.

3. The Explanatory Adequacy of a Representationist Construal of the Paradigms of A Priori *Truth*

In the second section of chapter 2, I claimed that a proper construal of *a priori* truth must explain, either in itself or as part of a larger theory, two kinds of characteristics of its subject matter: those which are possessed by any kind of truth, and those which are specific features of the *a priori* instances under scrutiny. Having said this, I compiled a list of the most important of these

explananda. The list included seven of the former and four of the latter type of phenomena. Since it was by no means meant to be complete, I argued that the joint explanation of these characteristics should be regarded as a minimal condition of adequacy for a construal of *a priori* truth.

In the previous section, I showed that the non-referentialist account advocated in this work can meet the third explanatory requirement specified in chapter 2. In particular, I argued that we can develop an (empirically confirmable) proposal of how our purportedly *a priori* claims about abstract domains may acquire both their representational truth conditions and their non-spatiotemporal intended referents. In this section, I shall show that the account can meet the other ten explanatory requirements as well. If my proposals prove to be correct, then we can conclude that the account under scrutiny provides an adequate characterisation of truth in the semantics of our paradigm *a priori* discourses, at least in view of the conditions of adequacy specified in chapter 2.

Fit with a General Construal of Truth

The first *desideratum* on our list toward a construal of the purported paradigms of *a priori* truth, which, arguably, must be maintained *vis-à-vis* a theory of any specific kind of truth, was that the account should harmonise with our general conception of truth. In particular, the semantic property characterised by the account must fall into the extension of our general concept of truth.

In chapter 3, I showed how the observation of this adequacy condition may lead someone with a broadly Tarskian, referentialist concept of truth to the denial of any non-referentialist response to Benacerraf's dilemma in philosophy of mathematics and the philosophy of our discourses about causally inert domains in general. The standard referentialist objection to the non-referentialist perspective is that an account that does not understand truth in terms of intended referents cannot be taken as a conception of truth. According to this perspective, what we

mean by this term is something essentially related to the subject matter of the relevant truthbearers. If truth is the property of having a subject matter (i.e. being about something) that obtains in the actual world, then our claims about abstract states of affairs cannot possess this property in virtue of the obtaining of some conditions in our head.

An advocate of the suggested representationist construal of truth in the semantics of our purportedly *a priori* discourses, of course, may argue that what our cognitive and linguistic practice in pure logic and mathematics (together with some assumptions of human knowledge acquisition) suggests is exactly that the received referentialist notion of truth is inadequate. According to this line of thought, what we mean by applying the term 'truth' in our discourses about abstract domains cannot be suitably specified in terms of the intended abstract subject matters.

Note that a referentialist may agree that the way we use the term 'truth' in our purportedly *a priori* discourses about abstract domains would disconfirm the referentialist reconstruction of our concept of truth if we could show that the term is applied in the same genuine sense in the latter discourses as in our empirical discourses about the natural world. Still, she may insist that we have no reason to suppose that the latter condition actually obtains. In fact, the very same evidence that makes her opponent abandon the standard referentialist construal of truth in the semantics of the relevant discourses can be invoked to query the idea that the term 'truth' is applied in the same genuine sense in these discourses as in the others.

The fact, however, that we assert something about truth in the semantics of certain discourses that we deny about it in the semantics of some others does not exclude that these claims are about the same semantic property. After all, most of our beliefs are not definitive of the nature of their subject matters. What is, on the other hand, clearly required from someone who insists that our cognitive and linguistic practice in discourses about abstract domains undermines the general referentialist construal of truth is an account of what makes the intended referent of our

notion of truth in the latter discourses the same as that in our discourses about the natural world.

Now, as has been briefly adverted to in chapter 3, from the perspective of a substantive realist, broadly use-theoretic semantics (a framework that emerged from our discussion of Horwich's minimalism in chapter 4), the opponents of standard referentialism may suggest that truth is best understood as *the property of possessing declarative use conditions that actually obtain in the world, in so far as the bearer of this property is a truth-apt mental or physical representation*. This understanding does not assume that truth has anything to do with the actual obtaining or absence of what its bearers purport to be about, and thus it is compatible with the suggested non-referentialist idea that our paradigm *a priori* claims about abstract domains are true or false in virtue of the actual obtaining of some representational conditions in our head. In so far as the construal that a non-referentialist provides is a construal of the correct declarative use conditions of our paradigm *a priori* beliefs, her account qualifies as an account of the paradigms of *a priori* truth, and thus satisfies the first adequacy condition set for such theories in chapter 2.

Objectivity of Truth

The second *explanandum* on our list for a construal of the purported paradigms of *a priori* truth (as well as, again, for a construal of any specific kind of truth) was the apparent objectivity of this semantic property, or in other terms, the fact that no one is ever conceptually prevented from committing epistemic mistakes (i.e. judge something to be true (or false) that is in fact false (or true)).

In chapter 4, I argued that the most natural way to explain this characteristic is to maintain that the truth value of our truth-apt representations is determined by the obtaining or absence of some conditions in the actual world, independently of anyone's actual knowledge or opinion of this circumstance. Putting it briefly, my suggestion was that the truth conditions of these representations are to be construed in a substantive realist way.

Apparently, the representationist construal advocated here of the truth conditions of our paradigm *a priori* claims about abstract states of affairs satisfies this explanatory requirement. One may object that it does not, since it implies that the obtaining or absence of the relevant truth conditions is the result of our epistemic activity, some stipulations that we make while developing our representations of the natural world. In response to this charge, however, an advocate of the suggested representationist account may observe that the reality of a certain condition requires merely that its obtaining or absence be independent of our actual knowledge, opinion or thought of this particular circumstance, rather than of our epistemic states and activities in general.[29] As has been emphasised earlier in this work, although the obtaining or absence of the crucial natural links among our representations is, indeed, the result of our epistemic activity in the spatiotemporal world, our knowledge, opinions or thoughts of these particular facts in our head are definitely not among the relevant epistemic determinants.

Summing up, since the suggested representationist construal offers a substantive realist understanding of the truth conditions of our paradigm *a priori* beliefs about abstract domains, an advocate of this account can explain the objectivity of *a priori* truth by reference to the fact that this property characterises its bearers in virtue of the obtaining of some substantive, real conditions.

Knowledge / Reliability of Evidence

The next phenomenon that we said an account of the paradigms of *a priori* truth, and in fact an account of any knowable kind of truth, must (arguably, in conjunction with our actual theory of human cognition) somehow explain is the possibility of

[29] Note that our intuition about truth in pure logic and mathematics is not that it obtains independently of any epistemic activity in the world, but instead that it obtains independently of what anyone ever actually thinks of this particular circumstance in the world.

knowledge acquisition, or reliable belief formation, of the relevant sort of truths.

As we have seen in chapter 3, the essential point behind Benacerraf's epistemological challenge to the standard platonist construal of mathematical truth (or its modified generalisation against the corresponding realist and referentialist construals of truth in discourses about causally inert domains) was that the account cannot satisfy the current adequacy condition, since it undermines the possibility of a reasonable explanation of how we can acquire knowledge or reliable evidence of what can be truly held about the relevant intended domains. In chapter 6, I argued that this charge against platonist theories of knowable truths is legitimate, and that a suitable construal of the paradigms of *a priori* truth must be naturalist and non-referentialist in character.

Obviously, the representationist account advocated in this chapter satisfies this requirement as well. It claims that the truth conditions of our paradigm *a priori* beliefs about abstract domains obtain (if they do) in our head, which means that our knowledge of or reliable evidence for what can be truly believed about these domains can be explained by reference to some causal mechanisms, whose nature is (in principle) just as open to systematic empirical study and characterisation as the nature of our knowledge or reliable belief formation of the natural world.

Intersubjectivity of Semantic Content

The next *explanandum* on our list in chapter 2 for an account of the paradigms of *a priori* truth, as well as for an account of any communicable kind of truth, was the intersubjectivity of the semantic content of the bearers of this semantic property. Although there are well-known sceptical considerations querying the existence of shared meanings, our successful daily communicative practice in pure logic and mathematics as well as in our empirical discourses about the natural world suggests that different people are still capable of entertaining (largely) the same ideas (i.e. thoughts that are about the same things and are applicable, in any given representational context, under the same

circumstances). So, a theory that cannot support an account of this phenomenon can hardly qualify as an adequate characterisation of the relevant type of truth.

In the first section of chapter 6, I conceded that the standard platonist construal of the semantic content (i.e. the truth conditions and the subject matter) of our claims about abstract domains can meet this explanatory requirement. If John's and Peter's respective thoughts *that one plus one plus one equals three* are supposed to be about, and also made true by the obtaining of, the same platonic conditions in the world, then any explanation of how John's and Peter's relevant mental symbols acquire their semantic contents (i.e. any account of the third *explanandum* specified in chapter 2) will *eo ipso* explain how these two subjects become able to share the above mathematical thought.

Note, however, that assuming the actual obtaining of the intended states of affairs cannot be a prerequisite for a successful account of our capacity to share and communicate ideas in different discourses. If it were, then we could have no such account of this capacity in the case of our beliefs about fictive or real-but-actually-uninstantiated entities.

The representationist account defended in this chapter preserves the idea that two people can entertain thoughts about the same abstract conditions (e.g. that John's and Peter's respective thoughts *that one plus one plus one equals three* are about the same abstract conditions), but it denies that the truth conditions of these thoughts are to be understood in terms of those abstract intended referents. Instead of adopting this referentialist construal, the account rather assumes that the truth conditions of our purportedly *a priori* thoughts about abstract domains obtain in our heads. What the advocates of this non-referentialist alternative suggest is that the fact whose reliable detection gives rise and justifies John's knowledge *that one plus one plus one equals three* is different from the one whose reliable detection is the grounding source of Peter's knowledge *that one plus one plus one equals three*. The former is supposed to be found in John's head, while the latter in Peter's. Now, the obvious question that an opponent of this account may raise is how this

difference between the respective truth conditions could ever be reconciled with the claim that the two representations under consideration have the same semantic content.

The non-referentialist answer to this question is that the truth conditions of a particular subject's (purportedly *a priori*) ideas of abstract domains obtain (if they do) not merely in this subject's head, but also in every single person's head who has acquired the ability to entertain those ideas in her own mind. Due to this construal, the semantic content of John's and Peter's respective thoughts *that one plus one plus one equals three* proves to possess the same content, despite the fact that the particular representational conditions whose obtaining gives rise to John's belief are clearly different from those whose obtaining inform Peter's belief. On the representationist account under scrutiny, both ideas are about the same abstract entities, and both are associated with the same (functionally identified) representational declarative use conditions in (the relevant) human heads.

One may wonder what makes it the case that John's and Peter's thoughts in the previous example are about the same abstract entities if there are no abstract entities in the world. As we said, the same question can be raised concerning our apparently co-referential thoughts about fictive or actually uninstantiated entities. The proper answer to this question can be derived from what has been said, in the previous section as well as in the first section of chapter 5, about our various capacities of developing new concepts from earlier acquired ones by subjecting the latter to some concept-forming operations, which alter the subject matter of their input concepts in a systematic way.

In the case of our particular example, John's and Peter's respective thoughts *that one plus one plus one equals three* are about the same abstract entities, because the corresponding conceptual elements of these thoughts also share their subject matter, while the latter condition obtains, because John and Peter developed these notions by the same kind of operations from corresponding concepts that represented the same obtaining conditions in the spatiotemporal world.

Similar response can be given to someone who wants to see why the relevant representational conditions in John's and Peter's heads can be regarded as fulfilling the same functional role. In the previous section, it has been argued that the conditions in question are established by the two subjects while they are developing their representations of certain properties and operations instantiated in the spatiotemporal world. Since in the case of the corresponding conceptual elements of the above thoughts in John's and Peter's heads these original natural properties and operations are supposed to be the same, it seems legitimate to maintain that the relevant representational features in the two heads have the same function in the natural world: they constitute the factual basis of analytic relations among representations, in different heads, of the same aspects of reality.

Summing up, similarly to her referentialist opponent, an advocate of the representationist construal articulated in this chapter can explain the intersubjectivity of the semantic content of our paradigm *a priori* beliefs about abstract domains by invoking the shared causal origin of the relevant semantically equivalent representations. An important novelty of the previous explanation is, however, that it does not presuppose that the truth conditions of these beliefs can be specified in terms of the intended abstract referents.

Observable Convergence of Beliefs

The next feature that occurred on our list in chapter 2 as an *explanandum* for an account of any type of truth was the observable convergence (or divergence) of our opinions concerning the distribution of the characterised semantic property. In our paradigm *a priori* discourses about abstract domains, the measure of convergence among various subjects' semantically equivalent beliefs is prominently high. Accordingly, an account of truth in the semantics of these discourses must support a suitable explanation of this high measure of convergence.

In the first section of chapter 6, it has been argued that if we had a belief-forming mechanism by which we could reliably detect the obtaining or absence of platonic truth conditions, then by reference to this capacity the advocates of the standard platonist construal of the paradigms of *a priori* truth could explain the observable convergence of our opinions concerning the distribution of this semantic property.[30] They could simply argue that different subjects' reliably generated beliefs about the obtaining or absence of the very same conditions must largely coincide.

As in the case of the previous *explanandum*, however, it must be noted that assuming the existence of a reliable information-conveying mechanism between minds and intended referential domains cannot be a prerequisite for a successful account of the observable convergence of our opinions in different discourses. If it were, then we could hardly develop such an account in the case of our converging beliefs about fictive domains.

The proponents of a representationist construal of the paradigms of *a priori* truth do not invoke the existence of such a mechanism in their account of this *explanandum* either. What they assume is rather the existence of a corresponding reliable information-conveying mechanism between minds and the domains of relevant truth conditions. In the case of our paradigm *a priori* beliefs about abstract states of affairs, the latter domain is, in their view, to be found in human heads. Accordingly, what they invoke in their account of the observable convergence of these beliefs is the existence of a reliable information-conveying mechanism between human minds and the alleged (obtaining or absent) representational truth conditions in human heads.[31]

[30] The primary purpose of chapter 6 was, of course, to show that we have no suitable ground to suppose the existence of such reliable belief forming mechanism.

[31] More precisely, the mechanism invoked must be "first-personal" between the subject's mind and the relevant representational conditions in her *own* head. The representationist construal under scrutiny does not exclude the possibility of empirical knowledge of the obtaining of these conditions in human heads.

Prima facie it may occur that a representationist cannot make use of the previous assumption, since, in her view, the particular conditions whose obtaining or absence is supposed to give rise, respectively, to John's and Peter's purportedly *a priori* beliefs about a certain abstract state of affairs are not the same: the conditions detected by John obtain in John's head, while those detected by Peter obtain in Peter's. Clearly, if these conditions were the same, then John's and Peter's reliable epistemic access to their obtaining or absence could explain the convergence of these subject's relevant co-referential beliefs. In absence of this identity, on the other hand, there seems to be no guarantee that the invoked reliable mechanisms give rise to the observed convergence.

Note, however, that what the explanation of this phenomenon seems to require, beyond the existence of the above reliable information-conveying mechanisms in our head, is not the identity, but merely the co-obtaining of the detected representational truth conditions. Returning to our previous example, if the representational conditions whose obtaining is meant to give rise to John's belief *that one plus one plus one equals three* obtain in John's head if and only if the corresponding representational conditions whose obtaining is supposed to inform Peter's belief *that one plus one plus one equals three* obtain in Peter's head, then the fact that both John and Peter develop their respective beliefs by reliable information-conveying mechanisms can explain the observable convergence of these beliefs.

So, what the proponents of the representationist account must show is merely that their theory implies the co-obtaining of the respective representational truth conditions of various subjects' purportedly *a priori* beliefs about a certain abstract state of affairs within these subjects' heads. Fortunately, the obtaining of this implication can be easily demonstrated. As we saw, the account assumes that the obtaining (or absence) of the relevant representational truth conditions of various subjects' purportedly *a priori* beliefs about a certain abstract state of affairs is constitutive of the semantic content of these beliefs. Accordingly, the fact that John and Peter are equally capable of entertaining

the idea *that one plus one plus one equals three* entails that the respective representational truth conditions of these beliefs in John's and Peter's heads co-obtain.

Summing up, the adoption of a representationist construal of the paradigms of *a priori* truth does not undermine the suitable explanation of the observable convergence of our paradigm *a priori* beliefs. Instead of assuming the existence of a reliable information-conveying mechanism between human minds and the intended abstract referential domains of these beliefs, an advocate of this construal can account for the current *explanandum* by invoking the existence of a similar cognitive mechanism which provides the subjects with reliable information of the obtaining or absence of the suggested representational conditions within their own heads. Since the account assumes that the truth conditions of these beliefs co-obtain in the heads of those subjects who can entertain the relevant paradigm *a priori* thoughts, the high reliability of the suggested belief-forming mechanism can explain the observable strong convergence of these beliefs.

Infinity of Semantically Non-Equivalent Truth-Apt Representations

The last feature on our list in chapter 2 that must be explained by an account of virtually any type of truth was that the property under scrutiny can characterise an infinite number of semantically non-equivalent bearers. In arithmetic, for instance, we can in principle entertain any particular member of the infinite thought-series *1+1=2, 2+1=3, 3+1=4*, etc. (i.e. any instance of the thought-scheme $n+1=k$, where k and n are natural numbers and k is the successor of n). Consequently, a proper construal of arithmetical truth must support an account of what makes it the case that there are infinitely many semantically non-equivalent representations that may be the bearer of arithmetical truth.

In chapter 6, it was noted that the standard platonist construal of the examined paradigms of *a priori* truth can meet this explanatory requirement by endorsing a simple referentialist construal of the semantic content of our paradigm *a priori* claims

and observing that the intended referential domain of these truth-apt representations is infinite in character. *Prima facie*, a similar explanation should be available for the advocates of the suggested representationist construal of the paradigms of *a priori* truth as well. After all, a representationist queries neither the referential aspects of the semantic content nor the infinity of the intended referential domain of our paradigm *a priori* beliefs.

Opponents, however, may object that the substantive realist (use-theoretic) semantical framework put forward earlier in this work undermines the adequacy of this standard explanatory strategy in the representationist's case. According to that framework, semantic contents are to be specified in terms of correct declarative use conditions. In the case of our paradigm *a priori* claims about abstract domains, representationists hold that these conditions obtain, if they do, within our heads. Supposing that there is only a finite number of these conditions, the representationist construal seems to entail that there can be only a finite number of semantically distinguishable purportedly *a priori* beliefs about abstract domains.

In response to this objection, a representationist may recall that claims with identical truth conditions may differ in semantic contents, because declarative use conditions are in general more fine-grained than truth conditions. As it was noted in chapter 4, the truth conditions of a truth-apt representation can be identified with the declarative use conditions of this complex symbol in an unembedded state. Clearly, the fact that two representations can be applied under the same circumstances in an unembedded state does not entail that they preserve their equivalence in larger representational contexts as well.[32] In view of this relation between truth conditions and semantic contents, a representationist may argue that her finitist conception of the domain of the relevant truth conditions is fully compatible with the idea that there are infinitely many semantically different, purportedly *a priori* claims about abstract domains.

---

[32] See fn. 15 in chapter 4.

A second objection to the representationist's view that the obtaining of finitely many conditions in our heads determines the truth value of all (determinately true or determinately false) paradigm *a priori* claims focuses on the representationist assumption that the conditions in question are natural relations among some representational constituents of the relevant claims. This time, opponents may argue that there seem to be infinitely many semantically basic constituents that we can combine into paradigm *a priori* claims about abstract domains. Moreover, it seems that we can conceive various infinite series of these truth-apt representations, in which no two members can possess the same analytically related constituents. If the above representationist assumption is true, then the existence of such series implies that our paradigm *a priori* claims must have infinitely many truth conditions.

Consider, for instance, the series of truth-apt arithmetic representations mentioned above. According to the representationist construal, the truth conditions of the idea *that one plus one equals two* are some links among the semantically basic constituents of this representation in our heads, and the same is supposed to hold of the truth conditions of every other member of the series. Now, it appears that no two members have exactly the same analytically related constituents. This implies that none of them can, according to the suggested representationist assumption, possess exactly the same truth conditions. Since the series consists of infinitely many arithmetic claims, it seems that a representationist cannot consistently hold that the truth value of these members is determined by the obtaining of a finite number of conditions.

In response, a representationist may query two premises in the previous line of thought. First, she may observe that in the actual world we can possess only a finite number of representations in our heads. Consequently, in the metaphysically thick sense, there are not infinitely many (mental or physical) representations in the world. The conceivability of the above infinite series of arithmetic thoughts requires, in a representationist framework, merely the conceivability, not the

actual existence, of an infinite number of distinct representational truth conditions, something that clearly obtains in the actual world.

The second representationist reaction concerns the opponent's assumption that no two members of the arithmetical series under scrutiny can have exactly the same analytically related constituents. *Prima facie* the assumption seems true. If we focus on the mental or physical symbols applied in these representations, then we find that each member of the series differs in at least one crucial constituent from any others. The first doubt concerning this claim emerges when we recognise that each member of the series is an instance of the same general arithmetical thought, the idea *that the sum of the number one and an arbitrary natural number equals the successor of that number*. Since we can establish the truth value of this general thought with our finite cognitive capacities, a representationist may reasonably assume that the truth conditions to be detected in this case obtain within the actual finite domain of representations in our heads. Further, since the truth of this general thought guarantees the truth of its infinitely many instances, she can also assume that the former conditions constitute the truth conditions of these instances as well.[33]

[33] One may object that the identity of the truth conditions of a general analytic thought with those of its instances would entail the logical equivalence of these representations, which is incompatible with the logical laws that govern universal quantification. My view is that in the case of analytic generalisations (i.e. generalisations that are analytically true), the equivalence holds whenever the analytic link guaranteeing the truth of the general idea obtains among its constituents other than the universal quantifier. For instance, I hold that the generalisation *that bachelors are men* is logically equivalent with, rather than merely entails, its own instances (e.g. the idea *that the tallest bachelor is a man*, the idea *that the second tallest bachelor is a man* etc.). Due to the analytic character of these truths, indeed, we cannot conceive any of them holding without the holding of the others. Of course, the representations are semantically different, because they do not share all their constituents. Nevertheless, the presence of their specific components does not affect their inferential relations, since these relations are fixed by the analytic link between the relevant common components of these representations. (Thanks to Daniel Isaacson for calling my attention to this objection.)

The opponent may insist that this point illuminates merely the inadequacy of the representationist assumption that the truth conditions of our paradigm *a priori* claims are natural relations among some constituents of these claims. In other terms, the opponent may grant that each member of the above series is made true by the obtaining of the same representational conditions, insist that no two members have exactly the same analytically related constituents, and therefore conclude that the suggested representationist characterisation of the particular truth conditions of our paradigm *a priori* claims cannot be adequate. In order to save this part of her conception, a representationist must show that our paradigm *a priori* thoughts may have analytically related common constituents even if no such constituents can be identified on the level of those mental or physical symbols that appear in the actual formulation of these thoughts.

The crucial observation that a representationist may invoke at this point is that the way we conceive an infinite set of truth-apt representations which has no two members with the same analytically related symbolic constituents is always by conceiving an infinite number of instances of a finite number of more general thoughts. The instances of a general thought, however, usually possess as semantic constituents most constituents of the general thought. Of course, at least one of these constituents must occur in the instances in infinitely many specified forms. This is what distinguishes the instances from the general thought. Nevertheless, these specifications do not annul the presence of the specified constituent in the particular instances. It merely terminates the explicit symbolic representation of this common semantic component in those thoughts. It is due to this fact that we cannot completely identify the common constituents of a certain multiplicity of thoughts merely by comparing those mental symbols that occur in the actual formulation of these representations.

In the case of our infinite arithmetic series, for instance, a representationist may argue that, despite appearances, most semantic constituents of the general thought *that the sum of the number one and an arbitrary natural number equals the successor of that*

number (viz. our concepts of addition, number one, natural number, successor, sameness and equality) appear in each of the conceivable instances of this thought, and that, corresponding to the suggested representationist assumption, it is the actual obtaining of a finite number of suitable natural relations among some of these common constituents that guarantees the non-referential truth of the relevant infinitely many arithmetic beliefs.

The illusion that there are no such common constituents emerges merely because some of these elements have no explicit symbolic representation in the chosen conceptualisations of the purported arithmetic contents. Our notion of natural number, for instance, does not appear explicitly in the thought *that one plus three equals four*. Note, however, that our concept of number three is a symbolic shortcut for the symbolically more complex representation of the successor of the successor of the only natural number that is not a successor. The two representations are associated with the same arithmetic content. Accordingly, the idea *that one plus three equals four* must involve, at least as an implicit semantic constituent, our notion of natural number as well. The only constituent of the above general thought that is clearly absent in its particular instances is the idea of universality (the notion of everything-of-a-certain-kind or everything *simpliciter*). This component is replaced by various individuating concepts, which turn our general notion of natural number into ideas of particular natural numbers. As we saw, our notion of number three can be regarded as a specification of our general concept of natural number.[34] The most important observation from the representationist's perspective is that the semantic relations of the applied individuating concepts need not be constitutive of the truth conditions of the resulting infinitely many arithmetic claims.

Summing up, an advocate of the suggested representationist construal can explain the conceivability of an infinite number of

[34] The claim that certain concepts can be regarded as specifications of some others is meant to have no implication concerning the actual genesis of the related symbolic elements. New concepts can be developed from earlier acquired ones by specification as well as by abstraction.

semantically non-equivalent paradigm *a priori* claims even if her account implies that the truth value of these claims is determined by the obtaining of a finite number of conditions in our heads. Following her standard referentialist opponent, a representationist may argue that the previous phenomenon is a consequence of two semantic facts. The first is that claims with different intended referents have different semantic contents, while the second that the intended referential domain of our paradigm *a priori* claims is infinite in character. The finitist implications of the construal concerning the domain of relevant truth conditions turned out to be compatible with this account on three different considerations: first, the representationist's general use-theoretic notion of truth conditions implies that semantic contents are more fine-grained than truth conditions; second, the finitist implications emerge only in so far as the construal is meant to characterise the truth conditions of our paradigm *a priori* representations in the actual world, where (in the metaphysically thick sense) there is only a finite number of such representations; and third, even those infinitely many paradigm *a priori* representations that we can at least conceive (as opposed to separately entertain) in this world are instances (and logical consequences) of a finite number of more general representations, whose truth value can arguably be determined by the obtaining of a finite number of conditions in our heads.

Apriority of Evidence

Beyond the above general features, whose proper explanation is arguably a minimal adequacy condition for a construal of virtually any kind of truth, in chapter 2 I collected four further characteristics that constitute an *explanandum* for a suitable account of the paradigms of *a priori* truth in particular. The first of these characteristics was that our knowledge of these paradigms is based on *a priori* evidence, or, in other terms, that it can be justified without reliance on experience.

Of course, in a sense this explanatory task can be trivially accomplished by any construal of the subject that supports an

acceptable account of our knowledge within the relevant discourses. The fact that we regard certain sorts of thoughts as the paradigms of *a priori* representations implies that our knowledge of their truth value must be also *a priori* in character. The situation is, however, far from being so simple, since our notion of apriority is not entirely void of empirical content. In fact, as we saw, it is understood in terms of independence of experience. Now, even if we have no definite concept of experience, nevertheless we understand that, for instance, any knowledge of the external part of the spatiotemporal world is based on experience. If this is so, however, then a construal which locates the truth conditions of the examined representations in the external part of the natural world cannot be consistently regarded as a construal of the paradigms of *a priori* truth.

In chapter 6, I acknowledged that in possession of an acceptable account of knowledge acquisition or reliable belief formation of abstract domains, the advocates of the standard platonist construal of those paradigms could meet this explanatory requirement. If they stood in the contact theorist camp, they could simply observe that our access to the alleged platonic truth conditions cannot be empirical, since it requires a specific epistemic capacity (different from the use of our external or internal senses) which connects our minds with entities outside the natural world. In contrast, if they opted for a no-contact epistemology, they could argue that our knowledge of abstract domains cannot be empirical, since it requires no contact whatsoever between our minds and the relevant platonic truth conditions.

Now, can a proponent of the suggested representationist construal of the truth conditions of our purportedly *a priori* claims about abstract domains also consistently assume that her account of knowledge acquisition or reliable belief formation within the relevant discourses is apriorist in character? Well, at the beginning of the second section of this chapter, it was noted that the main motive behind the non-referentialist's adoption of the suggested representationist version of naturalism concerning the above

conditions is that this construal seems to observe our (slightly indefinite) idea of experience, which implies that all knowledge of the external part of the natural world and almost all knowledge of our own bodily states is empirical.

Of course, some construals of the *a priori* / *a posteriori* distinction would be incompatible with the representationist position. For instance, if one supposed that any knowledge of the spatiotemporal world is by definition empirical, then one could not consistently maintain that the truth conditions of our paradigm *a priori* claims obtain in the domain of representations within our heads. As it was noted in the first section of chapter 1, however, the conflict between the representationist conception and such an understanding is merely terminological. The substantive representationist point is that our claims about abstract domains (and our analytic claims in general) are true or false in virtue of the obtaining of some representational conditions in our heads, and our knowledge of these representational facts constitute a natural kind, which can be contrasted with our knowledge of other parts of the natural world. The representationist classifies this kind as *a priori*, because it is traditionally regarded as the paradigm of this type knowledge, but if her opponent wants to use this term with theoretical connotations that make radical empiricism a plausible position in epistemology, she can accept this alternative convention as well, and merely insist on the previous substantive part of her doctrine.

Reconciled to the representationist terminology, one may still wonder what a representationist could tell us about the essential feature of what she calls *a priori* knowledge. The simplest representationist response to this question is that the desired feature is exactly that the facts detected by this type of knowledge are representational, and they obtain (among others) within the thinker's own head. Our knowledge of them is *a priori*, because it can be acquired before getting any (empirical) information of what obtains in the rest of the actual world. Beyond this response, a representationist may add that the exact specification of the nature of *a priori* knowledge and justification is the task of our empirical sciences, and it must emerge from a painstaking

empirical inquiry into the work of human brains in the course of the relevant types of knowledge acquisition.

In possession of these responses, we can conclude that an advocate of the suggested representationist construal of truth in the semantics of our purportedly *a priori* discourses about abstract domains can adequately account for the fact that our knowledge within these discourses is based on *a priori* evidence by maintaining that a piece of evidence, or the justification that it provides for a belief that is based on its recognition, is *a priori* if and only if it is generated by a reliable causal mechanism within human heads that conveys information of the obtaining or absence of truth conditions within a subject's system of representation to her knowing mind.

Necessity of (the Paradigms of) A Priori *Truth*

The second specific *explanandum* for an account of the paradigms of *a priori* truth that appeared on our list in chapter 2 was that the relation of this semantic value (or its opposite) to its bearers is necessary in character. My claim that there are three chairs in front of me could be false even if it is actually true. In contrast, my idea that there are three prime numbers between 70 and 80 is not merely true, but it is necessarily true.

In chapter 6, I argued that an advocate of the standard platonist construal of the paradigms of *a priori* truth can provide a relatively simple explanation of this phenomenon. Namely, she can maintain that the necessity of the relevant truths (or falsities) is due to the necessity of the obtaining (or absence) of those platonic truth conditions that our paradigm *a priori* claims purport to be about.[35] The proponents of the suggested representationist

[35] In absence of a reasonable account of how we could actually detect the modal character of the obtaining of various truth conditions, however, this referentialist conception of the factual ground of modality creates a serious explanatory problem in the epistemology of our cognitive discourses in general. Note that Lewis's alternative strategy, which explains the modal status of necessary truths and falsities by reference to the simple (non-modal) obtaining of referential conditions in realistically construed non-actual worlds results in similar

construal cannot follow this route, since on their view the relevant truth conditions obtain (or not) contingently in some segments of the actual spatiotemporal world. But is there an alternative account of this *explanandum*? Can the simple (i.e. non-necessary) obtaining of some natural conditions in the actual world guarantee the necessary truth (or necessary falsity) of our paradigm *a priori* beliefs?

My claim is that the proper response to this question is also positive. To see why, we should merely recognise that the necessity of the truth (or falsity) of a certain thought can be also a consequence of the semantic fact that the obtaining (or absence) of the relevant truth conditions is constitutive of the semantic content of this thought. In section 2, I argued that the suggested representationist construal of the paradigms of *a priori* truth is compatible with such an account of this *explanandum*. After all, the analytic relations obtaining among our representations, the suggested representational truth conditions of our paradigm *a priori* claims, can be arguably regarded as constitutive of the semantic content of the related symbols as well as the claims composed of them. According to this understanding, the claim that there are three prime numbers between 70 and 80 cannot be false, because in absence of the (actually obtaining) non-referential truth conditions of this representation the declarative use conditions of at least some of its constituents would also differ from the actual ones.

Putting it briefly, the representationist construal advocated here can explain the observable contrast between contingent and necessary truths and falsities by invoking the difference between two kinds of truth conditions in the actual world: those whose simple obtaining (or absence) merely determines the truth value of some claims, and those whose simple obtaining (or absence),

epistemological difficulties in so far as it leaves no room for a viable account of how we could actually detect what obtains in the suggested realistically understood but causally closed alternative possible worlds. Lewis (1986).

beyond determining this value, also contributes to the semantic content of those claims.[36]

Applicability of A Priori *Knowledge in the Empirical Sciences*

The third specific feature to be explained by an appropriate account of the paradigms of *a priori* truth on our list in chapter 2 was the applicability of our knowledge of these truths in our empirical sciences (as well as in our ordinary theorising) of the natural world. To take the simplest example, if we learned from experience that Peter put two apples on an empty table, and then John added three others to these two, then we can know without counting the group again, merely relying on our *a priori* knowledge that two plus three equals five, that there must be five apples lying on the table now.[37]

[36] Note that this account can explain the existing ambiguity in the modal character of those truths and falsities as well, whose subject matter is the actual obtaining (or absence) of the invoked representational truth conditions in our heads. Consider, for instance, the claim that the representational truth conditions of the idea that there are three prime numbers between 70 and 80 obtain in my head. According to the suggested representationist construal, this claim is presumably true. But is this truth necessary or contingent? Well, we may easily conceive a world in which the representational conditions in question do not obtain. Note, however, that in such a world those missing conditions would not constitute the truth conditions of this mathematical claim. In fact, in such a world the claim could not be made at all. (In a purely syntactical sense, the symbols applied could exist there as well, but the semantic contents associated with these entities would be different from the actual ones.) What this means is that there is a sense in which the above truth is necessary, and an other in which it is contingent in character. We can easily conceive a world in which the truth conditions of an actually thinkable analytic truth would not obtain, but we cannot conceive a world in which this content could be entertained in thought without being true. (Thanks to Hartry Field for turning my attention to this consequence of the representationist construal advocated in this work.)

[37] The history of science provides more complex illustrations of this phenomenon. The greatest insights behind the early modern emergence and later development of our current scientific conception of the world are famously associated with the invention of illuminative thought experiments, *a priori* reasonings that seem to provide new synthetic knowledge of some laws in nature. The best known examples of such applications of *a priori* (logical and

The crucial question to be answered in order to meet this explanatory requirement is what sort of relation between the truth conditions of the relevant claims makes it the case that our purportedly *a priori* knowledge of abstract domains can help us discover what is necessarily true about the spatiotemporal world. In chapter 6, I observed that, in view of the stipulated metaphysical gap between the intended platonic and natural realms, an advocate of the standard platonist construal of the truth conditions of our paradigm *a priori* beliefs seems to have no satisfactory response to this question, and, consequently, her theory would remain inadequate even if it proved to be compatible with an acceptable account of our knowledge of the relevant abstract domains.[38]

In contrast, an advocate of the suggested representationist framework can provide a relatively simple answer to the previous question. As we saw, in this framework the truth conditions of our paradigm *a priori* claims about abstract domains are relational states of affairs obtaining among our representations that we

mathematical) knowledge in natural science can be found in Stevin (1605), Galilei (1632), Galilei (1638), Newton (1687), Maxwell (1871), Einstein (1905), Einstein, Podolski and Rosen (1935), and Schrödinger (1935). The most influential classical theories of the nature of thought experiments include Mach (1897), Duhem (1906), Koyré (1939), Popper (1959), and Kuhn (1977b). For more recent literature on this topic, see Brown (1991), Horowitz and Massey (1991), Miščević (1992), Sorensen (1992), Hull, Forbes and Okruhlik (1992), Haggqvist (1996), McAllister (1996), Norton (1996), Gendler (2000), and Hitchcock (2004).

[38] One may try to defend the platonist construal in this regard by first maintaining that the alleged metaphysical gap does not exclude the existence of certain isomorphisms between what obtains in the contrasted platonic and natural realms, and then declaring that the existence of these isomorphisms would perfectly well explain the observable applicability of our purportedly *a priori* knowledge of platonic domains in our empirical sciences of the natural world. The problem with this response is that a platonist has no independent evidence of the obtaining of this enormous coincidence between the facts of the two realms. Her evidence for her pure logical and mathematical beliefs cannot be taken as such, because, as I attempted to show in chapter 6, we have no reason for assuming that our epistemic grounds in these paradigm *a priori* disciplines provide reliable information of what actually obtains in a platonic realm. To this measure, the platonist's failure in the two explanatory respects mentioned here is still related to each other.

develop in our head in the course of our cognitive engagement with our natural environment. In section 2, we saw also what a representationist can in principle say about the emergence of these conditions and their envisaged semantic links to the above claims. Most importantly, the account in question assumed that the conditions are analytic relations among our representations of abstract entities that we developed by abstraction in the third specified sense of the term in our head from some earlier acquired ones of various aspects of the natural world. So, the question to be answered by a representationist boils down to this: what makes it the case that our knowledge of the obtaining or absence of these analytic links can help us discover what is (necessarily) the case in the spatiotemporal world, or rather, in non-epistemic terms, why do these analytic relations, the obtaining truth conditions of our true representations of abstract domains, reflect so well what actually obtains in the spatiotemporal world?

The representationist answer to this question is the following. If the account advanced in section 2 is correct, then the truth conditions of our paradigm *a priori* claims about abstract domains are identical with the truth conditions of their various "applications" to the natural world. The truth conditions of our arithmetic claim that two plus three equals five, for instance, are identical (among others) with those of our "applied arithmetic" claim that Peter's two and John's three apples must constitute a group of five apples on the table. The identity in question is ensured by the way we develop our ideas of the relevant abstract entities from our ideas of some properties that may obtain, and thus be observed, in the spatiotemporal world.

An important consequence of this semantic fact is that our *a priori* acquired logical and mathematical knowledge cannot fail to help us discover or clarify, in principle, what (necessarily) obtains in the natural world, unless the relevant "applications" fail to be true about this world too. The latter circumstance, however, could obtain only if the analytic relations determining the truth value of these "applications" were introduced by us among symbols with unrelated (fine-grained) referential contents (i.e.

unrelated synthetic declarative use conditions).[39] A careless act like this would turn (the relevant part of) our system of representation inconsistent. For instance, if we stipulated that water is whatever has the chemical structure CO_2, then the claim that water is partly constituted of carbon would become analytically (and thus necessarily) true. On the other hand, unless we radically changed the actual referential content of our notion of water (or that of our notion of carbon dioxide), the same claim would be clearly false about the actual world. Note, however, that an analytically true representation cannot be consistently false about the actual world (if the referents of its constituents exist in that world). Apparently, maintaining consistency requires that we create analytic links between symbols only if the referential contents of these entities stand in a suitable (part-whole or identity) relation with each other.

In section 2, we saw that there are two major types of situation in which we actually introduce an analytic link between two elements of our system of representation. Sometimes, the link emerges when we define a new concept to represent a certain aspect of the world in a simpler way than its *definiens* did before. In the actual world, this is how we develop our concept of bachelor. In such cases, the referential contents of the related symbols cannot fail to stand in the required relation with each other. The synthetic declarative use conditions of our notion of bachelor, for instance, have presumably included from the outset as a proper part those of our concept of man.

At other times, we seem to introduce the link among symbols with independently established referential contents. This is what we seemed to do when we reduced our idea of water to our chemical concept of H_2O, and maybe also when we first defined our notion of being a double in terms of our notion of being a single and our idea of addition. In these cases, the required relation of the relevant referential contents does not

[39] I insert the term 'fine-grained' here to recall that the notion of referential content that I rely on in this part is not the coarse-grained Fregean one, which was shown to lead us to the collapsing conclusion of the slingshot in chapter 4.

seem to be automatically guaranteed. As we observed in section 2, however, the link that we actually introduce when we "reduce" a non-analytically acquired notion to some others is (1) strictly speaking not among these semantically independent representations, and (2) even loosely speaking among these entities only if the intended referential contents in question sufficiently overlap (as in the case of "ordinary" water and H_2O) or coincide (as in the case of being a single together with another single, on the one hand, and being a double, on the other). In fact, what happens in the case of these theoretical reductions is that we replace the relevant semantically independent notions with some others, which are analytically connected and possess suitably related referential contents that are also either identical or just slightly different from those of the respective antecedents.[40]

Due to this aspect of our concept-developing practice, the referential contents of our analytically linked notions never fail to meet the semantical requirement specified above. Accordingly, our analytic claims never fail to be necessarily true about their intended referential domains.[41] This implies that the analytic "applications" of our paradigm *a priori* claims of abstract domains to the natural world are also always true about this world, which in turn explains why our purportedly *a priori* knowledge of the former domains can help us discover what is necessarily true about the latter realm.

[40] It is this identity or sufficient overlap between the initial and modified referential contents that "legitimises" or motivates the strictly speaking rather confusing linguistic practice that we express the newly introduced notions with the same physical symbols as their analytically unrelated antecedent (e.g. we keep the term 'water' to speak about the newly conceptualised stuff that is essentially constituted of H_2O).

[41] They may, of course, fail to apply to the world if applicability is understood in the stronger sense requiring the actual obtaining of the synthetic declarative use conditions (or the existence of the intended referents) of all basic constituents of the relevant analytic claims. Note, however, that our current *explanandum* is not the applicability of our *a priori* knowledge of abstract domains in this strong sense of the term, but its applicability when the referents of the symbolic elements of its "applications" do exist in the natural world.

Summing up, an advocate of the suggested representationist construal of the paradigms of *a priori* truth can explain the applicability of our knowledge of these truths in our empirical sciences of the natural world by invoking two facts about the truth conditions of the relevant *a priori* claims: first, that they are identical with those of the conceivable "applications" of these representations to the natural world; second, that their obtaining is a consequence of our stipulative work, whose nature guarantees that the "applications" in question are never false about the natural world.

Abstractness and Infinity of Intended Domains

The final specific characteristic that occurred on our list in chapter 2 as an *explanandum* for an account of the paradigms of *a priori* truth was that the intended referential domains of the bearers of these paradigms are abstract and often infinite in character.[42] The question to be answered in this case is how the bearers of these paradigms can determinately refer to entities within such domains.

Earlier I argued that an advocate of the standard referentialist construal of the subject can provide a suitable answer to this question and successfully explain our capacity to think of or speak about abstract and infinite domains. Nonetheless, as we saw, the referentialist conception cannot be regarded as adequate, since it can explain the objectivity of its subject matter only if it involves a platonist understanding of this entity, a construal that supports no reasonable account of how we can acquire this type of knowledge, and why we can successfully apply it in our empirical sciences.

Fortunately, the adoption of an alternative non-referentialist conception of the same subject affects merely our idea of the

[42] Note that the possible infinity of referential domains, in itself, is not a specific *explanandum* for an account of the paradigms of *a priori* truth. After all, our capacity to develop ideas of infinite universes is manifest in our broadly physicalistic discourses as well.

relation of the truth conditions of the relevant bearers to their intended referents. What an advocate of the standard construal took to be plain identity, a proponent of the suggested representationist construal regards as a less intimate relation. In particular, she maintains that the truth conditions of our paradigm *a priori* representations are identical with some analytic relations that may obtain among the basic constituents of these contentful entities.

In accounting for the current *explanandum*, the phenomenon that our paradigm *a priori* beliefs can determinately refer to entities in abstract and often infinite domains, a representationist can simply follow her referentialist opponent, and argue that the referential content of the basic constituents of these beliefs is fixed by those cognitive mechanisms that underlie the actual development of these representations in human heads. In so far as at the bottom they are all developed from earlier acquired notions (with determinate spatiotemporal referents) by abstraction in the third specified sense of the term, their fixed referential content will be, quite understandably, non-spatiotemporal in character.[43] On the other hand, since in possession of a limited number of basic concepts (again, with determinate referents) of a certain domain we can recursively develop, at least in principle, an unlimited number of new (determinately referring) representations of previously not represented entities of that domain, the possible infinity of what we can think of or speak about in the relevant *a priori* discourses should not strike us as a surprise either.

With this result, I have finished the examination of the adequacy of the suggested referentialist construal of the paradigms of *a priori* truth. If the accounts advanced in this section are correct, then we can conclude that in contrast to its referentialist alternatives discussed in chapters 4, 5 and 6, the non-referentialist conception proposed in this work can satisfy all major

[43] In chapter 5, I argued that determinate reference to spatiotemporal entities can be explained within a "metaphysical realist" perspective as well.

explanatory requirements set for an account of this subject in section 2, and thus qualifies as a minimally adequate conception of what we took to be the paradigms of *a priori* truth. In the last section of this chapter, I shall return to the broader perspective of our current investigation and argue that the representationist conception advocated here can be taken as a suitable defining characterisation of the nature of *a priori* truth in general, and in possession of this definition, we can derive a corresponding characterisation of the nature of *a priori* knowledge, justification and evidence as well.

4. A Representationist Construal of A Priori *Knowledge and Truth*

I started this work with the announcement that I would argue for a particular naturalistic characterisation of *a priori* truth and knowledge, a conception that preserves the categorical distinction between *a priori* and empirical beliefs, largely observes the received application of the contrasted terms, and, together with our best empirical theories of the world and human cognition, explains the most important characteristics of our purportedly *a priori* claims and beliefs. In order to ensure the greatest initial agreement concerning the premises of our investigation, I proposed to start with the examination of truth within what we usually take to be the paradigms of *a priori* discourses (i.e. within pure logic and mathematics), develop an adequate account of this subject, and then consider whether the resulting conception would suggest us a suitable real definition (i.e. a characterisation of the nature) of *a priori* truth and knowledge in general.[44]

With the previous section of this chapter, the first part of the above plan has been completed. After presenting a list of the most important conditions of adequacy that the envisaged

[44] The project thus conceived corresponds to standard reductive analyses in the empirical sciences, such as the earlier discussed scientific reduction of (our concept of) water to (that of) H_2O.

account of the paradigms of *a priori* truth must (arguably without exception) satisfy, I argued for two major tenets concerning the subject under scrutiny. First, I attempted to establish that the truth conditions of our paradigm *a priori* beliefs cannot be understood along the received referentialist lines. Second, I argued that a particular representationist construal of these conditions can pass the above test, and thus qualify as a minimally adequate characterisation of the paradigms of *a priori* truth. In possession of these results, we can turn now to the second part of our project, and consider whether any of these findings could be invoked in a minimally adequate characterisation of the nature of *a priori* truth, knowledge, evidence and justification in general.[45]

First, it may be worth noting that the two features mentioned in the previous paragraph are independent of each other. In other terms, the fact that the truth conditions of certain beliefs are representational does not entail that they are non-referential, and *vice versa*. As it was mentioned earlier, the truth conditions of our beliefs *about* the obtaining of the alleged representational truth conditions of our paradigm *a priori* claims are arguably also representational but referential. On the other hand, one may argue that the most plausible construal of the truth conditions of our truth-apt normative claims (e.g. in ethics or epistemology) is non-referentialist but not representationist in character. So, what would follow concerning the extension of our notion of *a priori* truth, if we invoked either or both of these features in the envisaged characterisation of the nature of this property?

[45] As it was mentioned in chapter 1, in this work apriority is meant to be a property that is primarily attributed to justifications. Derivatively, however, it can be also attributed to pieces of knowledge, beliefs, propositions, judgements, sentences, utterances and truths due to the relevant epistemological features of these entities. Here I suppose that an *a priori* truth is a truth that can be known *a priori*, and a truth can be known *a priori* if and only if our belief in it can be justified *a priori*, without reliance on experience. For more on this topic, see section 1 in chapter 1.

Well, consider first what would happen if we stipulated that both non-referentiality and representationality are essential for certain truths to qualify as knowable *a priori*. Such a construal would entail that neither our claims about the obtaining of the alleged representational truth conditions of our paradigm *a priori* beliefs, nor our claims about the occurrence of causally inert normative values in the spatiotemporal world, could be taken as *a priori* in character. On the other hand, as was argued in section 2, our standard analytic claims about the spatiotemporal world do possess the above two characteristics, and therefore, according to this understanding, they would also fall into the category of (truth-apt) *a priori* representations.

As a second option, we could stipulate also that the essential feature of an *a priori* knowable truth is that it is non-referential. This would render our standard analytic claims as well as, arguably, our claims about the occurrence of causally inert normative values in the spatiotemporal world *a priori*, while our claims about the obtaining of the alleged representational truth conditions of our paradigm *a priori* beliefs empirical in character.

Finally, we could say that a certain truth is *a priori* if and only if it is representational. This would imply that beyond our standard analytic beliefs our claims about the obtaining of the representational truth conditions of these beliefs are *a priori* too.[46] On the other hand, this understanding would be incompatible with the idea that our knowledge of the occurrence of causally inert normative values in the spatiotemporal world, if possible at all, is *a priori* as well.[47]

[46] Whether or not we acknowledge that our claims about the obtaining of these representational conditions can be analytic in character, it is relatively easy to see that they can be taken as synthetic claims about a certain segment of the actual spatiotemporal world. To say that we can acquire *a priori* knowledge of the obtaining of these conditions is nothing else than to recognise the existence of synthetic *a priori* claims about the actual world.

[47] A further notable consequence of all these construals would be that our synthetic thoughts about abstract domains could not be regarded as *a priori* either. Of course, this does not mean that we must classify them as empirical. This would follow only if every thought would be either *a priori* or empirical. But why should we assume this? After all, was not our earlier conclusion of these truths that they cannot be known by us anyhow?

Note that neither of these definitions would preserve the purely epistemological character of our distinctions between *a priori* and empirical truth, knowledge, evidence and justification: while the contrast between referential and non-referential conditions is clearly a semantic one, that between representational and non-representational conditions is one (with semantic connotations) within our actual ontology. So, whichever option we choose, our notion of apriority will lose its traditional epistemological significance.[48]

Nonetheless, I believe that our notion of non-referentiality and notion of representationality are not equally appropriate candidates for becoming an *analysans* of our notion of *a priori* knowledge and truth. In so far as we regard apriority as a property that primarily characterises ways of justifications, the semantical question whether or not certain truth-apt representations are *about* their actual truth conditions seems largely irrelevant. The fact, for instance, that our beliefs about the obtaining of the (allegedly representational) truth conditions of our (pure) logical and mathematical claims are about their own truth conditions does not seem to affect at all the way in which we actually justify these beliefs. In contrast, the envisaged location of these conditions within the spatiotemporal world, and especially the relation of this location to that of knowing minds seems much more significant from the perspective of our current investigation.

For these reasons, I believe that from among the three analytic proposals specified above the intuitively best strategy for us is to opt for the third (i.e. to stipulate that a certain truth is *a priori* if and only if it consists in the obtaining of some representational conditions in a subject's head). As we observed, this understanding implies that our knowledge of causally inert

[48] As it was mentioned in chapter 1, a contrast is supposed to be purely epistemological if it can be characterised by reference to epistemologically significant properties, such as fallibility, transparency or fundamentality, alone.

normative properties or values must be either analytic or empirical. A proper assessment of this implication is clearly beyond the scope of the current work. Here and now it suffices to say that the observable characteristics of our linguistic and cognitive practice in normative discourses, such as ethics and epistemology, do not seem to stand in obvious conflict with the above assumption.

In possession of this representationist account of *a priori* truth, we can formulate now the corresponding representationist construals of *a priori* knowledge, evidence and justification. According to these construals, a piece of knowledge is *a priori* if and only if it is justified *a priori*, by invoking *a priori* evidence, and a piece of evidence, or the justification that it provides for a belief that is based on its recognition, is *a priori* if and only if it is generated by a reliable causal mechanism within human heads that conveys information of the obtaining or absence of truth conditions within a subject's system of representation to her knowing mind.[49]

Despite their largely "ontological" (or maybe semantical) character, these construals nevertheless preserve an important connotation of the traditional epistemological notion of apriority. Namely, they preserve the idea that, in a sense, *a priori* knowledge is prior to any further knowledge of the actual world. The priority in question was already briefly indicated in section 2. There we observed that our synthetic knowledge acquisition of the world is realised by two consecutive epistemic mechanisms. First, we develop an access to our representations, and recall those circumstances under which they can be correctly applied in the particular declarative representational context under scrutiny. Second, we discover whether or not the recalled circumstances actually obtain in the represented world. By adopting the suggested representationist construal of apriority, we can say now that the first sort of mechanism is what underlies our *a priori* knowledge acquisition of various domains. Accordingly, we may

[49] Note that the resulting notions enable us to distinguish in a relatively sharp way between *a priori* and introspective knowledge, justification and evidence as well.

conclude that on a representationist understanding of apriority, some *a priori* knowledge is constitutive of any propositional knowledge of the actual world.

With this conclusion, we can close now the second part of our investigation as well. After arguing that truth within our paradigm *a priori* discourses is best understood in a non-referentialist and representationist manner, in this section I suggested that by reference to the latter characteristic we can provide a minimally adequate specification of the nature of *a priori* truth, knowledge, justification and evidence in general.[50]

Summary

In this chapter, I attempted to show that, contrary to its referentialist alternatives, a specific naturalist version of non-referentialist realism about truth in the semantics of our paradigm *a priori* discourses satisfies all major adequacy conditions set for such a theory in chapter 2. I called this version a representationist account of the relevant truths, since it assumes that the truth conditions of our purportedly *a priori* claims about abstract domains obtain (if they do) in the realm of representations within our heads, rather than in the domain of the represented abstract states of affairs.

In section 1, I argued that, in view of our actual cognitive and linguistic practice, we have no reason to suppose that a non-referentialist construal of certain truths is conceptually objectionable.

In section 2, I started the demonstration of the adequacy of the representationist construal by providing an (in principle empirically confirmable) ontologically naturalist account of how our purportedly *a priori* claims about abstract domains may

[50] As it was emphasised in chapter 2, one can challenge the adequacy of the representationist construal advocated here by identifying some further characteristic of *a priori* truth and knowledge that the account under scrutiny can explain neither in itself, nor as part of a larger theory of the world.

acquire their determinate semantic relations to their (non-referential) representational truth conditions, on the one hand, and their non-spatiotemporal intended referents, on the other. In possession of this account, I concluded that the construal under scrutiny can satisfy the third explanatory requirement set for such a theory in chapter 2.

In section 3, I confronted the proposed construal with the ten other explanatory requirements collected in chapter 2, and argued that this version of non-referentialism is compatible with an appropriate account of all those phenomena whose joint explanation we agreed upon to regard as a minimal condition of adequacy for a conception of the relevant truths, and thus it constitutes also a suitable response to Benacerraf's original or modified and generalised dilemma in the philosophy of our purportedly *a priori* discourses about abstract domains.

Having shown that truth within our paradigm *a priori* discourses is best understood in a non-referentialist and representationalist way, finally, in section 4, I returned to the initial broader perspective of our investigation, and argued that by reference to the latter characteristic we can provide a minimally adequate specification of the nature of *a priori* truth, knowledge, justification and evidence in general.