Int Ontology Metaphysics
DOI 10.1007/s12133-011-0083-5

A Priori and A Posteriori: A Bootstrapping Relationship

Tuomas E. Tahko

Abstract The distinction between a priori and a posteriori knowledge has been the subject of an enormous amount of discussion, but the literature is biased against recognizing the intimate relationship between these forms of knowledge. For instance, it seems to be almost impossible to find a sample of *pure* a priori or a posteriori knowledge. In this paper, it will be suggested that distinguishing between a priori and a posteriori is more problematic than is often suggested, and that a priori and a posteriori resources are in fact used in parallel. We will define this relationship between a priori and a posteriori knowledge as the *bootstrapping relationship.* As we will see, this relationship gives us reasons to seek for an altogether novel definition of a priori and a posteriori knowledge. Specifically, we will have to analyse the relationship between a priori *knowledge* and a priori *reasoning*, and it will be suggested that the latter serves as a more promising starting point for the analysis of aprioricity. We will also analyse a number of examples from the natural sciences and consider the role of a priori reasoning in these examples. The focus of this paper is the analysis of the concepts of a priori and a posteriori knowledge rather than the epistemic domain of a posteriori and a priori justification.

1 Introduction

The topic of this paper is the a priori/a posteriori distinction, and especially the relationship between these forms of knowledge. Usually the distinction is expressed in terms of knowledge that precedes experience and experiential knowledge or non-empirical and empirical knowledge. We will use this wide conception as our starting point, although our preferred analysis of the a priori is in fact in terms of metaphysical modality rather than experience.[1] In the course of the paper, we will

[1]See Tahko (2008) for a definition of a priori knowledge in these terms as well as a preliminary account of the bootstrapping relationship.

T. E. Tahko (✉)
Department of Philosophy, History, Culture and Art Studies, University of Helsinki, Helsinki, Finland
e-mail: tuomas.tahko@helsinki.fi

Published online: 24 August 2011

see why this wide interpretation of the a priori/a posteriori distinction in terms of experience might not be sustainable. It should be noted that we will not discuss a priori *justification*: in Albert Casullo's terms, we are more interested in a *nonreductive* rather than a *reductive* approach to a priori knowledge—the first is concerned with the analysis of the concept of a priori knowledge, while the latter is concerned with a priori justification (cf. Casullo 2003: 10). Although 'a priori' and 'a posteriori' are arguably epistemic notions, our primary focus is not epistemological but metaphysical. In fact, we have more sympathy towards the idea that a purely epistemological account of these notions is not even possible or not particularly insightful at any rate. Leaving that aside, it should at least be made clear that the project at hand does not fall in the same category as much of recent discussion about a priori knowledge and its relationship with a posteriori knowledge (cf. BonJour 1998; Boghossian and Peacocke 2000; Casullo 2003; Peacocke 2004). What we are interested in is on how scientific and philosophical knowledge, assuming that they involve both a posteriori and a priori knowledge, are structured, and how knowledge accumulates. Another important aspect of the discussion concerns the relationship between a priori *knowledge* and a priori *reasoning*. It will be suggested that because it seems to be exceedingly difficult to define the former in a satisfactory manner, we should turn our attention to the latter; it will be suggested that one feasible definition of a priori knowledge is to be had in terms of a priori reasoning.

In the second section, we will examine whether the usual way to distinguish a priori and a posteriori knowledge is plausible, and argue that it is not. We will see that a sharp distinction between these forms of knowledge is not only very difficult to put forward but also largely unnecessary. However, we still maintain that we are dealing with two distinct forms of knowledge, even if they are fundamentally connected. These results are important in themselves, but they also underline the need for a new definition of a priori knowledge.

The primary purpose of this paper is to sketch a theory about the interaction of a priori and a posteriori knowledge which will also explain why it is so difficult to fully distinguish them. In the third section, we will introduce the idea that a priori and a posteriori knowledge are in a constant bootstrapping relationship. A priori knowledge generally advances in very small steps: we introduce an a priori proposition, which we then attempt to verify by a posteriori means; this is the core of the bootstrapping relationship.

The fourth section goes into the details of the bootstrapping relationship and develops the idea that there is an established a posteriori framework which can be augmented by the process of bootstrapping—depending on the result of our attempts to verify a given a priori proposition, it can either become a part of the established a posteriori framework or it can be discarded because it is incompatible with the a posteriori framework. This is how knowledge accumulates, and we wish to demonstrate that the structure of our current a posteriori framework reflects the bootstrapping relationship between a priori and a posteriori knowledge. It should be emphasised that justification does not enter the picture: the a posteriori framework is fallible and we may wish to change it according to different input. Different beliefs might be justified at different times, but this is not the story that we are engaged with here.

In fifth section, we will examine the 'roots' of the bootstrapping relationship, that is, whether bootstrapping begins with a posteriori or a priori knowledge and, more

generally, whether there are good reasons to think that one or the other form of inquiry is more fundamental. Here we will consider the law of non-contradiction as a possible example of a fundamental piece of a priori knowledge that might serve to ground the bootstrapping relationship.

The sixth section concludes the paper with an analysis of the implications of the proposed understanding of the relationship between a priori and a posteriori knowledge.

2 How Can We Distinguish A Priori and A Posteriori Knowledge?

There are many well-known problems concerning the a priori/a posteriori distinction, but these rarely receive a detailed analysis in contemporary literature. For instance, Laurence BonJour (1998: 7–11) mentions two of the most apparent problems concerning the a priori and experience: the problem of how we define 'experience' and how the a priori is supposed to be 'independent' of it. In the first case, the problem is to determine the correct scope of experience. Do mental processes count as experience? How about mathematical or philosophical reasoning that relies on certain learned patters? Should only perceptual information count as experiential? How does memory fit in with all of this? The second apparent problem involves issues concerning concept acquisition as a precondition for a priori knowledge, but also independence of experience, namely whether a proposition is indefeasible by experiential information. Our approach to these problems is slightly different from that of BonJour's, as we are not interested in giving an account of a priori justification. But the requirement to clarify where the a priori stands in regard to experience is equally pressing for anyone interested in the notion.

For us, it will not be worthwhile to attempt to settle these questions by traditional means; in fact, we believe that they perhaps cannot be settled by traditional means. We suspect that this is because there is a more fundamental problem underlying the a priori/a posteriori distinction—one that undermines even the obvious solution of restricting a posteriori knowledge to perceptual information and a priori knowledge to non-perceptual information. The problem we have in mind is the following: no matter how strictly one might restrict knowledge to experience, hence classifying a certain proposition as a posteriori knowledge, it seems that there are always also some a priori elements present.[2] Consider some examples of perceptual a posteriori information that Colin McGinn takes to be particularly easy:

> [T]he scattering of birds causes you, via the belief that birds have scattered, to infer, with the help of a number of other beliefs, that there is a cat in the vicinity; the deaths of individual men (inter alia) may cause you to form the belief that all men are mortal; and so on (McGinn 1975–76: 199).

[2] See also Horvath (2009: 199–201), who suggests that coming to know the metaphysical necessity of propositions will always include a priori elements, even if there may be some a posteriori elements as well. This once again highlights the problem because *both* a priori and a posteriori elements would seem to be present.

On the face of it, these examples may indeed seem like clear cases of a posteriori knowledge. However, it seems that the process by which we might infer that there is a cat in the vicinity from the perceptual information concerning the scattering of birds and perhaps other things is not quite as simple as one might think. The crucial word here is 'infer'. Inference is a form of reasoning; it means that we can deduce something from a set of premises. The question is: what is the nature of this form of reasoning? Can it be accurately described as 'a posteriori reasoning', or is that even a sensible notion? Would we not be inclined to say that all forms of reasoning are a priori? One could of course suggest that only pure perceptual information, whatever that may be, is truly a posteriori and anything that we might deduce from that perceptual information is in fact a priori. However, this is not a common understanding of aprioricity. The point of the example is to demonstrate that the line between aprioricity and aposterioricity seems rather arbitrary: presumably all complex beliefs require deduction, but the deductive process often starts from empirical information. These problems call for an analysis of a priori reasoning and its role in distinguishing a priori and a posteriori knowledge.[3] There is a tension here because on the one hand it seems that we must have some perceptual information as well as the relevant concepts to be able to deduce anything at all, but on the other hand the acquisition of any kind of knowledge whatsoever would appear to require some reasoning, which is at least arguably always a priori.[4]

It is here that the notion of justification is often thought to be of assistance, as it is plausible that the same belief could have been reached in virtue of different things. Accordingly, we might say that the belief that there is a cat in the vicinity was justified due to certain empirical factors and perhaps previous knowledge about the behaviour of cats. But this does not appear to solve the problem: we may examine the inference pattern from the scattering of birds to there being a cat in the vicinity and discover that it is either fully independent of experience or contains some a posteriori elements, but unless we have a previous theory about the boundary between aprioricity and aposterioricity, then we have no means to settle how the proposition that there is a cat in the vicinity was justified. An obvious solution is to say that whenever there are empirical elements present, we are dealing with a posteriori knowledge, but because of the problems mentioned above, this would seem to imply that everything is a posteriori knowledge.[5] The line of thought that we wish to briefly entertain is the opposite: whenever there are a priori elements present, we are dealing with a priori knowledge. It should be emphasised though that this is a question that precedes the debate concerning justification, and it is even independent of it to some extent. We are looking for an analysis of the notions of aposterioricity and aprioricity, including their relationship, rather than a theory about how beliefs are justified. The questions are of course related, as once we have an analysis of these notions, and then we can proceed to define a priori and a posteriori justification accordingly.

[3] For more discussion, see Tahko (2008).

[4] Here's one reason why the notion of *a posteriori reasoning* does not seem to make much sense: our rational capabilities are plausibly independent of experience and the only a posteriori part involved with reasoning is pure perceptual information.

[5] If Philip Kitcher (1984) is right, this is even true of mathematics.

One thing that may support the controversial idea that we should adopt a wider sense of the a priori has even been voiced by McGinn himself, for he thinks that logic can be fairly straightforwardly classified as a priori in virtue of it being a purely deductive science (McGinn 1975–76: 199–200). Now, if deduction is a mark of the a priori, and we need deduction to be able to form a proposition concerning a certain state of affairs, such as there being a cat in the vicinity, then it would appear that any kind of propositional knowledge includes a priori elements. It may be that some of the underlying information originated in our senses, in fact this is surely the case, but once we have this preliminary, purely perceptual information, any further knowledge must be derived by a priori means. Perhaps this is not surprising, but if we adopt the hypothesis suggested above, namely that the presence of a priori elements dominates our analysis of the a priori/a posteriori distinction, then it seems that aprioricity pervades our knowledge. Yet, it would not appear to be appropriate to classify knowledge derived from perceptual information as a priori, regardless of the fact that we might have to use deduction in forming any propositions that concern perceptual information as well.

All this suggests that there might be something wrong with the usual way that we classify a priori and a posteriori knowledge because the distinction is not sensitive enough to account for the different elements underlying knowledge. That is, the obvious alternatives for the analysis of the notions of aprioricity and aposterioricity seem to result in nearly all knowledge falling to just one of the two categories. Furthermore, it is not clear how we should distinguish between a priori reasoning and a priori knowledge. If everything that is justified a priori is a priori knowledge, then it seems that propositions derived with the help of a priori reasoning constitute priori knowledge. But rarely if ever is a proposition derived strictly by a priori means, so does that make most knowledge a posteriori after all? One thing is clear: it is exceedingly difficult to pick out an example of a posteriori or a priori knowledge that would be pure. Accordingly, a dilemma emerges: we will have to either admit that the notions of a priori and a posteriori knowledge are fundamentally ambiguous because it seems impossible to pick out a clear example, or we have to redefine these notions altogether.

3 Introducing the Bootstrapping Relationship

There appears to be a demarcation problem concerning a priori and a posteriori knowledge. It is telling that a great deal of the literature concerning this distinction consists of attempts to settle the status of certain disciplines in regard to the a priori/a posteriori distinction—without much consensus. Perhaps this approach is fundamentally flawed: it may be impossible to determine a sharp boundary between the a priori and the a posteriori. In fact, if we consider how knowledge typically accumulates, the idea of classifying certain disciplines as a priori or a posteriori seems implausible. Surely we need both kinds of knowledge to engage in any philosophical or scientific activity, or so we will argue in this section. A priori reasoning, as was suggested in the previous section, is necessary for any kind of rational analysis, but a posteriori knowledge is also necessary if we wish to determine how things are in the actual world. We must have some kind of a story about what the role of a priori reasoning is here. In what follows, it will be suggested

that most if not all supposed cases of a priori knowledge are simply cases where the proposition at hand is reached primarily with the help of a certain deductive process which is a natural candidate for what we might wish to call a priori reasoning. This enables us to offer a preliminary definition of a priori knowledge. Perhaps even more importantly, many if not all propositions that are supposedly a posteriori are also reached (partly) with the help of the very same deductive process, that is, a priori reasoning. This suggests that there is an important connection between the a priori and the a posteriori. Our purpose in this section is to examine that connection and to introduce an alternative understanding of a priori reasoning.

It may be helpful to illustrate the connection between the a priori and the a posteriori with an example. Suppose that you get into your car to go to work in the morning, but for some reason the car does not start. Before doing anything else, you would probably wonder what different possible explanations there could be for this. Maybe the battery is empty, perhaps the car is out of petrol or possibly there is an electronic fault. The next thing to do is to check each scenario: does the battery work, is there petrol in the tank, are all the electrical systems functioning correctly? Of course, the person doing this could just as well be a car mechanic, and no doubt he/she could think of a lot longer list of possible reasons for the car not starting. In any case, the situation calls for a rational analysis of the possible reasons that could explain a certain observation, namely the car failing to start, as well as an empirical verification or falsification of these possible reasons. Now, if what was said in the previous section is correct, it is at least arguable that the rational analysis of the possible reasons for the car not starting is a priori by nature. The analysis does of course rely on previous a posteriori knowledge about cars and their functioning, but combining this knowledge with new empirical knowledge from the situation at hand and forming an hypothesis, an explanation for the car's behaviour, would appear to require more than just purely perceptual information—it requires deduction.

The point of the previous example is simple enough, although it would perhaps be much more plausible if we replaced cars and mechanics with, say, subatomic particles and physicists. A priori reasoning, however, is not just for scientists and philosophers, and this is exactly the point we wish to make with the previous example. In any case, all we need to take from this example is that when we acquire new empirical information we proceed to analyse it in terms of scenarios, possible explanations, which we consequently test by empirical means. In scientific contexts this is even more obvious: consider the reaction to the infamous double-slit experiment.[6] The strange empirical results observed in this experiment resulted in the wildest of scenarios of what could possibly explain the phenomenon, and to a certain extent, we are still trying to answer that question. The relationship between empirical observations and deductive reasoning concerning their nature is not always so straightforward though. Sometimes a very explanatory hypothesis may emerge without reference to any specific empirical observation; we will consider examples of this in what follows.

The implicit definition of a priori reasoning that one may extract from the above example has been discussed in this journal before (cf. Tahko 2008). For our current

[6] The double-slit experiment is an experiment in quantum mechanics, which demonstrates the wave-particle duality of light.

purposes, it is sufficient to note that the idea is that the modal element in our reasoning—scenarios of what is possible—is a sign of the a priori. In other words, a priori reasoning concerns possibility and accordingly a priori propositions are modal propositions. This contrasts with the point of Kripke (1980) that aprioricity does not imply necessity, which is usually considered to mean that the a priori cannot be analysed in modal terms. But we contend that with the correct interpretation of modality, this does not have to be the case.[7] We can analyse a priori knowledge in terms of metaphysical modality if we adopt the view that modality is grounded in essences, following Kit Fine (1994). However, it may be best to support the account with a number of examples rather than a detailed analysis of the metaphysics of modality. The main thesis of this paper, namely that the a priori and the a posteriori are in a bootstrapping relationship, is independent of this analysis in any case. All we need to keep in mind now is that there is a further story to be told about the modal content of the possible scenarios reached with the help of a priori reasoning.

The idea of bootstrapping can be recognized already in the example concerning the car refusing to start, but it will be helpful to consider more scientific examples. Let us take a genuine example from the history of science: consider the phenomenon of gravitational redshift.[8] The effect of gravity on light was already predicted by Newton's work, but Newton's results were inaccurate because of some false assumptions, namely Newton's corpuscular theory of light. If light is understood as an electromagnetic wave instead, then the phenomenon of gravitational redshift becomes puzzling, as it appears that the wavelength of light could only change from one place to another if the flow of time also changes. This mystery was of course solved by Einstein's special theory of relativity, which models how the flow of time can indeed change—although it was not until Einstein's theory of general relativity that he was able to produce the correct value for the redshift phenomenon. All this work was theoretical, but empirical verification soon followed: the Pound-Rebka experiment in 1959, which measured gravitational redshift, is usually considered to be the final verification.

We can already extract the philosophical content of this example. What we have here is a fine example of the bootstrapping relationship. Gravitational redshift is something that we would not have bothered to look for unless there were some non-empirical, theoretical reasons to do so. These reasons were already introduced by Newton's theory, which represents the first a priori step in our bootstrapping sequence. However, experiments with light showed that it behaves much like an electromagnetic wave, and the original theory concerning gravitational redshift could not accommodate this. Thus, we have an a posteriori falsification of the original theory. Einstein's work represents a new a priori step from the established a posteriori framework, and after his theory was empirically verified, it also became a part of the a posteriori framework. Here is this bootstrapping sequence as a diagram, starting from the bottom (Fig. 1):

[7] This also reflects the distinction between the modal status of a proposition and the truth of a proposition, see Horvath (2009) for more discussion.

[8] Gravitational redshift refers to the change in the wavelength of light and other electromagnetic radiation when it travels from a stronger gravitational field to a weaker one.

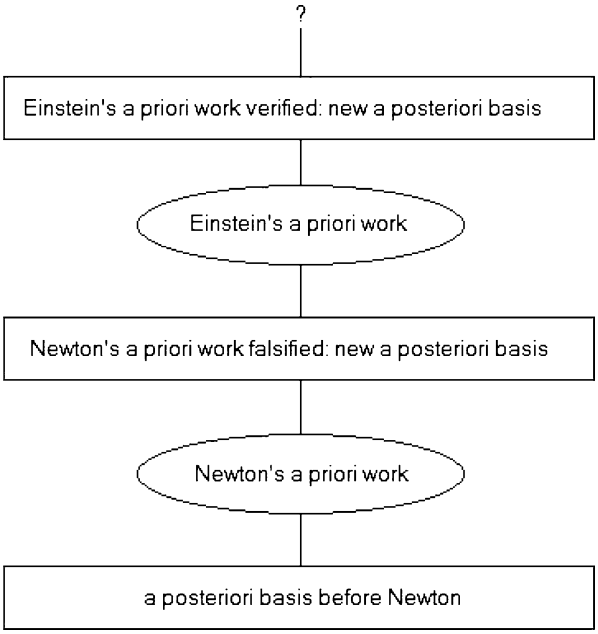

Fig. 1 Diagram showing a bootstrapping sequence

The diagram in (Fig. 1) is of course very simplified: only one line of a priori reasoning is illustrated, whereas generally there will be a vast number of different lines of a priori reasoning emerging from any given a posteriori basis. Some of them make it to the new a posteriori framework, some of them do not. Also, in this example we are dealing with whole theories rather than simple a priori propositions—this is primarily to illustrate how the idea can be applied to theory change as well. The basic idea behind the bootstrapping relationship is as simple as this, and it could certainly be accommodated in different theories about the a priori. However, it is possible to elaborate the idea somewhat, and for this we need to introduce some further material.

4 The Bootstrapping Relationship in Detail

We can imagine the bootstrapping relationship as a tree model: we start with a single trunk that represents the established a posteriori basis, a framework of empirical information. The branches of the tree represent lines of a priori reasoning, a priori propositions. But not all of these a priori propositions are true, some are merely possible. Similarly, there are a number of possible explanations for a faulty car not starting, but only one of them is the correct explanation. The dotted lines in Fig. 2 represent these possible but non-actual lines of explanation, they are dead ends. However, when we have a solid branch, one that represents a true a priori proposition, it can serve as a basis for a new tree that branches onwards.

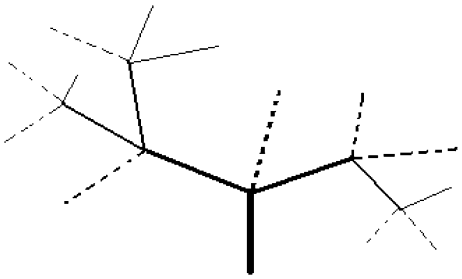

Fig. 2 The tree of knowledge

We might call the tree in Fig. 2 the tree of knowledge, albeit only a tiny section of it is illustrated here. At any given time, the tree of knowledge has a solid trunk, the established a posteriori basis, which we use as a starting point. For someone working in theoretical physics, this starting point includes our current knowledge of physics; for someone fixing car, it includes the report of what is wrong with the car and the knowledge of that particular car model. From this a posteriori basis, the physicist and the car mechanic can proceed by testing different scenarios, different a priori branches. Many of these branches turn out to be dead ends, although this may have to be determined by empirical means. Sometimes the scenarios are even incompatible with the established a posteriori basis, but in our tree of knowledge, these branches have been omitted, as it is probable that they are simply the result of human error. Eventually, a scenario that satisfies all the tests will emerge, a solid branch has been found. For the car mechanic, the story will end here: he/she knows why the car does not start and may proceed to fix the problem. For the physicist, one problem has been solved, but it is likely to be the source of a number of further problems, and the tree branches onwards.

The story so far leaves out a number of important details. Perhaps the most important of them is fallibility. It was already mentioned that branches based on fallacious a priori reasoning due to human error have been omitted in Fig. 2. Errors like this are usually spotted quickly though, because someone will surely spot the suspect line of a priori reasoning, even if not the person who came up with it in the first place; this is how the scientific community works. However, there is a more serious type of error that we are very familiar with. This is the type of error where we end up following one of the branches depicted with a dotted line in Fig. 2. How does this happen? Well, for the most part, the solidity of any given branch will be determined by a posteriori means—this is exactly why we have decided to call the relationship between a priori and a posteriori knowledge a bootstrapping relationship. When a scientist comes up with a new theory, presumably by a priori means, we proceed to test that theory with the help of empirical experiments; that is how we decide whether the theory is correct or not. But empirical information is fallible: results are open to a number of interpretations and they may be inaccurate, even the whole setting of the experiment may be flawed. What this means is that we can never be quite certain that the branch we are currently on is truly a solid branch rather than a dead end. Indeed, any branch that we have taken in the past may turn out to be rotten, no matter how certain we were about its solidity. We do not need to

look far for real life examples, Euclidean geometry may be one of the most notorious. In fact, backtracking to a previous branch and choosing another line altogether seems to be commonplace, we see this in all natural sciences. Accordingly, we may have to be more conservative when choosing a new branch, and it may even be worthwhile to pursue parallel branches, despite the fact that all but one of them will turn out to be a dead end. This is indeed a common practice in science: consider competing lines of research in theoretical physics, say, different theories of quantum gravity, such as string theory and loop quantum gravity. We know that only one of these theories can be correct—in fact they might all be incorrect.

The tree of knowledge is also deceptive. Although the bootstrapping relationship between a priori and a posteriori knowledge has proved to be very efficient and reliable in general, it is fundamentally fallible. We have no means to guarantee that we are on a solid branch, and to this extent, the established a posteriori framework is always subject to revision. Moreover, the tree is in a constant flux, and it may always turn out that a branch that was once abandoned turns out to be the correct one. Consider phlogiston, the oxygen-like substance that was commonly believed to exist in the seventeenth and eighteenth centuries, but turned out to be a very rotten branch indeed. The phlogiston theory was nevertheless consistent and it was pursued in great detail, all the empirical data available seemed to support it—until Antoine Lavoisier's discovery of oxygen, that is. Perhaps we are now in the position to say that the phlogiston theory will never re-emerge, there is just too much empirical data that speaks against it. That is not to say, however, that our current theory is necessarily the correct one, a change as radical as the one from phlogiston to Lavoisier's discovery of oxygen is certainly possible, even if it seems very unlikely.

It could be objected at this point that the process being described here is not one concerning a bootstrapping relationship between a priori and a posteriori knowledge at all, but rather the familiar story about scientific inquiry concerning scientific hypothesising and empirical verification of scientific theories. To a certain extent this is indeed the case, but we have been using scientific theories as an example exactly because they are a particularly good example of how the bootstrapping relationship works. It is not a very good objection to this project that it resembles the familiar story about how scientific knowledge accumulates—any plausible theory about the relationship between a posteriori and a priori inquiry will surely have to account for scientific knowledge, and if the suggested definition of a priori reasoning is correct, then this type of reasoning is the most crucial part of scientific inquiry. Moreover, the scenarios depicted by the branches of the tree of knowledge need not be just scientific scenarios; philosophical thought experiments and even counterfactual situations that we come up with in our daily lives all fall within the scope of the model. The purpose of this model is to illustrate how a priori reasoning functions, and the hypothesis is that a priori reasoning is required for all of these different rational activities: scientific hypothesising, philosophical thought experiments and thinking about counterfactual situations in everyday contexts.

One key question that remains is of course the interpretation of this model, namely, if the model correctly depicts the use of a priori reasoning, what is the use of a priori reasoning based on, what is the nature of the possible scenarios? Indeed, the

previous objection is perhaps based on the idea that the model will have to be interpreted in the lines of the familiar story about scientific hypothesising. The interpretation of the model is something that was already briefly discussed in the previous section.[9] The suggested interpretation is that these scenarios concern metaphysical modality. However, the model may be accurate even if this interpretation is not the correct one, and it will not be possible to defend this interpretation in any detail here. Our purpose is simply to present the model in its general form, and it is only an advantage if it can be combined with different accounts of aprioricity. In any case, it should be noted that the scenarios are not to be interpreted as conceptual possibilities; we also wish to strongly resist reducing a priori reasoning to conceivability or conceptual analysis. Why is this? Well, if the branches of the tree were simply conceptual possibilities, then the tree would quickly grow beyond what is manageable: everything from philosophical zombies to Disney fairy tales is conceivable; indeed, conceivability appears to be entirely unrestricted. But the tree of knowledge cannot be unrestricted, we have to be able to rule out metaphysically impossible scenarios that might nevertheless be conceivable, as otherwise we would never be able to make progress due to the infinite number of (conceptually) possible scenarios. Accordingly, because we clearly can make progress, there must be a method for delimiting the space of possible scenarios in such a way that metaphysically impossible branches can be dismissed outright. This supports the suggested interpretation: if a priori reasoning is a reliable guide to metaphysical modality, then our rational capabilities are automatically immune to the threat of metaphysically impossible scenarios. Keeping this in mind, let us turn to a final question concerning the foundations of the tree of knowledge.

5 Grounding the Bootstrapping Relationship

The question that we turn to now concerns the foundations of the tree of knowledge. No tree grows without roots and the same must presumably be true of the tree of knowledge. Metaphors aside, there must be a beginning to the bootstrapping relationship: if knowledge accumulates in the manner described, then it seems that we need a scenario, an a priori proposition, which we proceed to test to get the bootstrapping going in the first place. But which came first, the a posteriori basis or an a priori proposition? As we observed in the second section, it could be argued that only pure perceptual information is a posteriori, but pure perceptual information lacks the structure that the a posteriori basis needed to ground the tree of knowledge requires. We must categorize this information somehow and this categorization is arguably an a priori process. This may sound alarming: surely we should not assume that experiential information lacks determinate features which could themselves be sufficient for knowledge? The idea that perceptual knowledge is propositional is not unheard of, and this would suggest that it might be able to serve as the roots of the tree of knowledge. Indeed, we do not wish to assume some sort of Whorfian hypothesis, namely that experience owes its structure to a categorization by humans.

[9] A more detailed discussion is available in Tahko (2008).

We should not be fooled by the above line of reasoning: perceptual information may indeed be structured before we categorize it—in fact, animals are presumably able to observe this structure as well—but we are looking for the basis of a distinctly rational framework of knowledge: scientific and philosophical knowledge. In other words, we are dealing with knowledge that accumulates by theory-forming and rigorous tests for the truth of these theories. It is certainly plausible that no matter how structured pure experiential input may be, there is some (a priori) work to be done before a theory can be formed on the basis of that information. Accordingly, the basic tools for theory-forming are at least required before the tree of knowledge can grow.

So, what are the basic tools of theory-forming? It would appear that they must be based on some previous a priori principles, as theory-forming is certainly a rational activity, and we have observed that there seems to be a connection between a priori knowledge and rationality above. What could these a priori principles be? Well, perhaps our primary candidate for such a fundamental principle might be something like the law of non-contradiction (LNC). It may be that these a priori principles also include the fundamental categories of reality, whatever they are—at any rate, they are what our initial categorization of perceptual information should presumably be based on. But there is an apparent problem in determining the roots of the tree of knowledge, as we have no conclusive means to determine whether these candidate principles, such as LNC, are really fundamental a priori principles, or rather just some of the very first branches in the tree of knowledge. If the latter is the case, then it might turn out that the principles are incorrect, due to the fallible nature of the tree. However, the genuine roots of the tree may very well be necessarily true. Why? Because they must be principles of the most fundamental kind, principles without which the tree of knowledge would be unable to grow.

Consider the law of non-contradiction. It does indeed seem to be a good candidate for a fundamental principle, and perhaps it is also necessary, but it has been challenged nevertheless (e.g. Priest 2006). This is not the place to consider challenges to the law of non-contradiction,[10] but if it really is a fundamental principle, then even a possibility of a challenge is disconcerting. If our tree of knowledge is based on LNC (perhaps among other fundamental principles), and there is a way to challenge LNC, then it seems that an alternative tree based on thoroughly different principles would be possible. Why should we be worried about this? Well, if there are multiple trees of knowledge—trees with different roots which might have nothing whatsoever in common—then the whole of our accumulated knowledge can be questioned. It seems that we do not have any obvious ways to ensure that we are climbing the right tree of knowledge. But perhaps there is a way. After all, we are not climbing the tree blindfolded: we are in constant interaction with the world, we test our a priori propositions. Even though every test is in principle fallible, it is rather unbelievable that everything in our tree of knowledge could be undermined—this is essentially a form of the no miracles argument. Surely the internal coherence of our established a posteriori framework suggests that we have got at least the basics right.

[10] These challenges have been discussed in detail in Tahko (2009).

Furthermore, there are good reasons to think that only one tree of knowledge is possible, that is, there could not be an alternative tree based on some other fundamental principles. If LNC is indeed a fundamental principle, it would appear to be necessary for any tree of knowledge. This follows from the idea that LNC is a precondition for rationality; any kind of accumulation of knowledge, which is what the tree of knowledge is, will plausibly be based on a rational process of some sort. This is exactly the process that we have been examining. We cannot discuss the assumed connection between LNC and rationality here in any detail though. It should suffice to say that although some have questioned this connection, it is certainly a very widely accepted idea that LNC is necessary for rationality, and if this is the case, then LNC is clearly a fundamental principle for any kind of tree of knowledge. In practice, this means that there can only be one tree of knowledge, as otherwise the trees would have to share their roots. But only one tree can emerge from one set of roots, as a single tree can accommodate all the different branches that could emerge from that set of roots. In other words, the fundamental principles determine all the possible routes available to us, all a priori propositions. Once these fundamental propositions are fixed, it does not make sense to talk about alternative trees, at least if our assumption about the necessary connection between LNC and rationality is correct.

6 Conclusion

We have now outlined a way to account for the apparent connection between a priori and a posteriori knowledge while still maintaining their distinctness. Much of what has been said is implicitly or explicitly at odds with standard accounts of a priori and a posteriori knowledge. So much worse for the standard accounts: if we wish to maintain the intelligibility of the notions of a priori and a posteriori knowledge, a novel analysis of their relationship is needed.

One particularly controversial ramification of this account concerns empirical defeasibility. It may seem that the account implies that a priori knowledge always requires empirical verification and that it can also be falsified by empirical information. While the first supposed implication is correct to the extent that we need empirical information to verify that an a priori proposition is true of the actual world, the second supposed implication is misguided.

As to the idea that we need empirical information to verify the actuality of an a priori proposition, the explanation is simple: for instance, there are a number of possible geometries, such as classic Euclidean geometry and modern non-Euclidean geometries, but we are presumably most interested in the geometry exemplified by the actual physical reality. Alternative geometries may be possible, but they are not actual, they lack empirical verification. If we acknowledge Kripke's idea that aprioricity does not imply necessity, then this idea should not be very controversial, as surely we need empirical information to verify something that is contingent, such as which geometry is the actual one.

We can now also see why the account at hand does not imply empirical defeasibility. If alternative, non-actual geometries are a priori, that is, they were reached by a priori means in the manner described above, then surely they are still a

priori, regardless of the fact that they are not actual. In other words, merely possible a priori propositions, branches of the tree of knowledge that turn out not to correspond with actuality, are a priori nevertheless. This is presumably true of Euclidean geometry as well, insofar as it is internally consistent. Empirical information can only establish the truth or falsity of an a priori proposition in the actual world (and even this is fallible), but the proposition has its a priori status by necessity. If this is correct, then Euclidean geometry as well as the phlogiston theory and a great number of other such abandoned theories are still very much a priori, albeit merely possible.

Ultimately we need a novel definition of a priori knowledge, and the line suggested here is that this definition can be given in terms of metaphysical modality, which can be further reduced to essence (Fine 1994; Tahko 2008). However, the main goal of this paper was not to defend this definition, but rather to give a plausible account of the relationship between a priori and a posteriori knowledge. This goal has now been established, and in the process we have seen that the traditional, sharp distinction between the a priori and the a posteriori is not feasible. Until the metaphysical status and relationship of these forms of knowledge has been settled, we can make little progress in specific epistemic domains. Hopefully what has been said here can be the first step towards a deeper understanding of the relationship of aprioricity and aposterioricity

References

Boghossian, P. and Peacocke, C. (Eds.) (2000) *New Essays on the A Priori* (Oxford: Oxford University Press).

BonJour, L. (1998) *In Defense of Pure Reason: A Rationalist Account of A Priori Justification* (Cambridge: Cambridge University Press).

Casullo, A. (2003) *A Priori Justification* (Oxford: Oxford University Press).

Fine, K. (1994) 'Essence and modality' In Tomberlin, J. E. (Ed.) *Philosophical Perspectives 8* (Atascadero, CA: Ridgeview), pp. 1–16

Horvath, J. (2009) 'The Modal Argument for A Priori Justification', *Ratio*, Vol. 22, No. 2, pp. 191–205.

Kitcher, P. (1984) *The Nature of Mathematical Knowledge* (Oxford: Oxford University Press).

Kripke, S. (1980) *Naming and Necessity* (Cambridge, Mass.: Harvard University Press).

McGinn, C. (1975–76) '"A Priori" and "A Posteriori" Knowledge', *Proceedings of the Aristotelian Society*, New Series, Vol. 76, pp. 195–208.

Peacocke, C. (2004) *The Realm of Reason* (Oxford: Oxford University Press).

Priest, G. (2006) *In Contradiction: A Study of the Transconsistent*, 2nd expanded ed. (Oxford: Clarendon Press).

Tahko, T. E. (2008) 'A New Definition of A Priori Knowledge: In Search of a Modal Basis', *Metaphysica*, Vol. 9, No. 2, pp. 57–68.

Tahko, T. E. (2009) 'The Law of Non-Contradiction as a Metaphysical Principle', *The Australasian Journal of Logic*, Vol. 7, pp. 32–47.