A Priori Truths

Greg Restall*
The University of Melbourne
greg@consequently.org

Philosophers love *a priori knowledge*: we delight in truths that can be known from the comfort of our armchairs, without the need to venture out in the world for confirmation. This is due not to laziness, but to two different considerations. First, it seems that many philosophical issues aren't settled by our experience of the world — the nature of morality; the way concepts pick out objects; the structure of our experience of the world in which we find ourselves — these issues seem to be decided not on the basis of our experience, but in some manner by things *prior* to (or independently of) that experience. Second, even when we are deeply interested in how our experience lends credence to our claims about the world, the matter remains of the remainder: we learn more about how experience contributes to knowledge when we see what knowledge is available independent of that experience.

In this essay we will look at the topic of what can be known *a priori*. We will start with some examples of truths which we have thought to be uncontentiously *a priori* known.

Examples

From Logic: I can know that *if* every student has either passed an exam or completed an assignment, *then* either every student has passed an exam, or some student has completed an assignment. I can know this without concerning myself with the details of which students have passed an exam or have completed an assignment, for I can reason as follows: Let's suppose every student has either passed an exam or completed an assignment. We want to show that either every student has passed an exam or some student has completed an assignment. Suppose I choose a student. By hypothesis, he or she has either passed an exam or completed an assignment. If the assignment is completed, then I can conclude that some student has completed an assignment, which gives us what I wanted to show. If this piece of reasoning fails for *every* student, then we see that *every* student has completed an exam, which also gives us our desired conclusion.

This piece of argumentation suffices to prove what we wished to show. It used no details of students, assignments or exams. The form of the reasoning would work for any claim of the structure: if every F is either G or H, then either every F is G or some F is H. Valid deductive reasoning seems to give us *a priori* knowledge on the basis of logical structure.

* Thanks to John Shand for the invitation to write this essay, for supererogatory patience in waiting for its arrival, and for very useful comments on the initial draft, to Allen Hazen for very many discussions on the issues contained herein, and to my graduate students at the University of Melbourne, who have put up with many more discussions about *a priori* truth than is healthy.

From Semantics: It seems that I can know that all bachelors are unmarried, without going to the trouble of interviewing bachelors and checking their marital status. I know this because if I find out that someone is married, then this counts decisively against their being a bachelor. I know this because this is how I understand the terms 'married' and 'bachelor.' It seems that the meanings of the expressions I use governs *how* I treat the evidence I may find. I don't wait to learn of the marital status of a number of bachelors to conclude a generalisation on the basis of this evidence. I use this generalisation to govern the evidence I encounter.

From Mathematics. Take a case of mathematical reasoning. A salient example in the development of our understanding of the *a priori* and is relationship with logic and the analytic is *the intermediate value theorem*, which states that every continuous function f of real numbers, where the value $f(-1)$ that it takes on the input -1 is smaller than 0, and whose value $f(1)$ that it takes on the input 1 is greater than 0, has *some* input value x between -1 and 1 where the output $f(x)$ is 0. A continuous function which is *below* the origin line at -1 and *above* the origin line at 1, must have *crossed* the line at some point between -1 and 1. In other words, any continuous path (a path without jumps or breaks) starting on one side of a line and ending up on another side of that line must have at least one point at which it *crosses* the line.

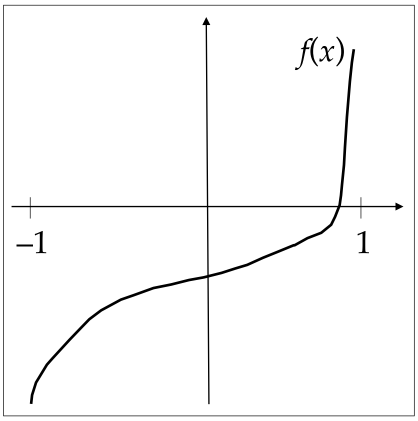

This is *obvious*. We can see that it is true in many cases, and furthermore, we may find it very difficult to know what it would be for it to be false in any case at all. For many years, mathematicians claimed that they could *see* that the theorem is true in full generality, without being able to offer anything resembling a proof. The key is the notion of *continuity*. A function is continuous if it has no breaks or jumps — and making this notion precise, in the development of the Calculus in the work of Bolzano and Weierstrass (see Coffa (1993), Chapters 2 and 3) was required for this piece of *a priori* knowledge to make the transition from something which seems obvious but is hard to justify, to something which could be proved on the basis of an analysis of the concepts involved.

From Indexicality: I may have amnesia and not remember where I am. However, wherever I am, if I say "I am here now" this seems to be something which is true, and which I can *know*. You use external evidence to find where I am, but I don't need to take in any evidence to convince myself that I am *here*. Evidence enters the picture when I want to know where "here" is, but as it stands, this is an item of *a priori* knowledge.

This case has interesting features, for while it seems to be a tautology for me to say that I am here now, that is not something that *you* can know *a priori*. You can know that whenever I say "I am here now" that I am speaking truly, but you cannot always know *what* I am saying. To see this consider me speaking to my spouse at the end-of-school-year concert, which I have unexpectedly been able to attend: "Zachary *doesn't know that*

I'm here now, he thinks that I'm still at work." I can know *a priori* that I am here, but Zachary, to learn that fact, must see me in the audience.

From Ethics: consider the judgement that making other people happy is, all things considered, a *good* thing to do. Perhaps when you consider this judgement you remember particular acts in which you have made others happy. But upon reflection, it seems that there's no reason to appeal to this or that experience, or this or that *evidence* that happiness is a good thing. Perhaps it is bound up in the very notions of happiness, others, and goodness, and we reasonably could come to that conclusion on the basis of reflection on those notions. Many believe that our fundamental ethical principles are known *prior* to the evidence. Such statements certainly don't look like they can be refuted (or proved) on purely empirically, and the fact that they concern what *should* happen seems to mean that mere description of how the world *is* cannot count decisively against the kind of claim about how it *should* be,

Are these good examples of *a priori* truths? How do they fare when we attend to the connections, if any, between *a priority*, necessity, analyticity, and infallibility? What have philosophers said about the notion of the *a priori*, and does the notion survive close scrutiny? These are the questions we will examine in the rest of this essay.

Definitions

As we have seen in our initial meeting with examples, an *a priori* truth is something that can be known independently of any particular evidence or experience. This rough and ready idea has been the basis of the claim to *a priority* for each of our examples. You do not need to know anything about the world in order to verify that if all Fs are G or H, then either all F are G or some F are H, or that all bachelors are unmarried, and I need to take in no evidence to know that I am here. Or so the reasoning has gone.

This does not mean that the notion of *a priori* knowledge is unproblematic. We have not given a characterisation what is necessary for knowledge to be independent of the external evidence, and neither have we had anything to say about what is to count as evidence, and of the evidence we have, what it means for some of it to be *external*.

Similarly, there is confusion on another aspect of the notion of independence. It is one thing to say that knowledge can be acquired without appeal to some antecedent experience. It is another to say that that item of knowledge cannot have its status as knowledge undermined by further experience of the outside world. This is a much stronger requirement. I may be convinced of some reasonably complicated mathematical result (such as the intermediate value theorem) by way of stepping carefully through some proof, with the help of a mathematically sophisticated friend. Although aided by props and outside support, this can count as *a priori* reasoning, for none of the props are essential to the content of what is conveyed, if I come to understand the proof. However, my confidence in my understanding of the proof could be shaky — although I think I have understood it, I may not be confident. In such a circumstance, my belief in the theorem may be undermined if my mathematically sophisticated friend then tells me that actually, that proof contained a subtle mistake and that mathematicians now believe that the theorem is false. My trust of my friend and in her authority as an expert may override my *a priori* knowledge, if I am not confident in that knowledge. This further information — that my expert friend tells me that the so-called theorem is *false* — is anything but a priori. This means that my knowledge of the intermediate value theorem, if it could be undermined by the claims of an expert, does not count as *a priori* in a stronger sense that requires unrevisability.

Whether the independence needed for a priority requires infallibility and unrevisability on the basis of any empirical evidence is a controversial issue (see, for example Casullo 2006, Kitcher 2000, Field 2000).

There are many other current debates concerning the nature of *a priority*. To take them in turn, we would do well to attend how thinking about the *a priori* has developed from Kant to the present day. While there is no doubt that the *a priori* played an important role in Ancient, medieval and early modern philosophy, the central importance of the notion, and its currently use is indelibly shaped by the work of Immanuel Kant.

The Synthetic A Priori

Before Kant's *Critique of Pure Reason,* the three allied notions of necessity, analyticity and a priority were not clearly distinguished. In Kant's defence of the synthetic *a priori*, the notions of the *analytic* and the *a priori* come apart. For Kant, an important class of truths can be known *a priori* but not through analysis. They are *synthetic a priori*. In our list of examples of *a priori* truths, the boundary between the analytic and the synthetic occurs in the split between logical and semantic/conceptual truths (which are analytic) and truths of arithmetic and geometry, which for Kant can be shown *a priori* only by the operations of the intuitions of time and space and not by analysis. While a considered discussion of Kant's account of the *a priori* is beyond us, we should venture into this territory just a little, for Kant's approach to the *a priori* set the scene for crucial developments into the 20th Century and beyond.

The distinction between the analytic *a priori* and the synthetic *a priori* is sharply drawn: for Kant, the boundary is found at the limits of what is possible through the *analysis* of concepts. Purely formal logic shows structural relationships between concepts (so Aristotle's syllogisms show formal and analytic conceptual relations between judgements), and the analysis of concepts into constituent parts grounds another kind of deduction between judgements. To say that *Gilles is a bachelor* is in part to say that *Gilles is unmarried*, and so, the conditional judgement *if Gilles is a bachelor, then Gilles is unmarried,* and the corresponding generalisation *all bachelors are unmarried*, can be shown to be true, *a priori*, by means of the analysis of concepts.

Nothing, according to Kant, can do the same thing with judgements such as simple arithmetic claims like: 8 + 5 = 13, or geometric claims, such as the claim that the interior angles of triangles add up to 180 degrees; let alone claims such as the *Intermediate Value Theorem*. In each case here, the reasoner must engage in some rational deduction in order to demonstrate the claims (and since the rational deduction is *pure* and not *empirical*, it is *a priori*, well enough), but his kind of reasoning goes beyond what is analytic. In the case of arithmetic reasoning, try as you like, you will not find the concept of 8 in the concept of 13, or *vice versa*. You will pass through 8 on the way to *counting to* 13, and you will hit 13 exactly after taking as many steps counting from 8 that one takes when counting to 5 — this is one way to show that 8 + 5 = 13, *a priori*. But, for Kant, this is not an analytic deliverance of some formal meaning: this is the kind of demonstration possible for one who has the concept of *time* — the concept of one thing coming *after* another, which is crucial to our practice of counting.

The same goes for geometric reasoning, but here, the requirement is not the concept of *time*, but the concept of *space*. We must use our spatial intuition in order to engage in *a priori* spatial reasoning. It is not for nothing that presentations Euclid's *Elements* are filled with diagrams, and geometrical reasoning is filled with instructions to "extend a line from point A through the intersection of lines *l* and *m* until it intersects line *k*" and the like — they are instructions to engage in our spatial reasoning by way of our spatial intuition. (Note here Kant does do not mean our *hunches* by the notion of "intuition" but rather, our grasp of concepts and the way we structure them together, pre-conceptually.) The findings of this synthetic *a priori* reasoning are as firm and as necessary as any analytic truth, despite the fact that it may be (in some sense) a contingent matter that we have the intuitions that we do. These intuitions of space and time are, for Kant, *pure*, because we have them antecedently to the acquisition of empirical concepts: we do not learn the pure intuitions of time and of space by example or by experiment: rather, we use these intuitions to structure the empirical intuitions we have *in* time and *in* space.

The 'contingency,' in some sense, of our having *these* temporal and spatial intuitions instead of others does not lead to the contingency of the judgements involving them: any more than the contingency of your having the concept of a *leg before wicket* and a *dismissal* means that it is not necessary that any batsman out by way of a *leg before wicket* has been *dismissed* by the *bowler*: what is contingent is your having the concept, and your ability to express (or understand) that necessary truth. It would not be any less necessary if we did not have the concepts needed to express it. Though it must be noted that in the case of the socially mediated practice of cricket, if no-one has the concept of *lbw* or *dismissal* or *bowler*, then it is plausible no-one can play cricket, either. Were we to not possess the spatial and temporal intuitions that we have, then we could not form these numerical or spatial judgements. Despite the internal and contingent nature of the forms of pure intuition, for Kant the link between the necessary and the a priori is fixed.

As we will see in the next section, philosophical orthodoxy has not remained Kantian on the matter of the *a priori*. Nonetheless, Kant has able defenders on the identification of necessity with *a priority* and the restriction of analyticity to a strict subset of the *a priori*.[*]

---

[*]Hanna (2001, 2006) is a clear and lucid contemporary exponent of this view.

The A Priori and the Analytic

It is perhaps not surprising that with Kant's key examples of synthetic a priori knowledge coming from mathematics, it was from developments in mathematics that revolutionary new ideas took their root. In the first instance, these ideas came from the nascent discipline of the calculus and the theory of real numbers and functions. We don't have space to recount the intellectual trajectory in any detail, suffice to say that by the 18th and 19th Centuries, since the pioneering work of Newton and Leibniz, the mathematical practice of differentiating and integrating functions was well established and largely well behaved, but — for both the pure mathematician and the philosopher — it was anything but well understood. What is it for a function to be continuous? Can we make the notion of "no jumps or gaps" precise, in a way that is amenable to reasoning? In the work of Bolzano and Weierstrass, Cauchy and Dedekind, through the 19th Century, advances were made on these fronts. Results such as the intermediate value theorem were provided with *proofs* in the 19th Century. The necessary ingredient wasn't the ingenuity to fill in a missing step in a calculation or a new technique for solving puzzles, but something much more fundamental, a new *definition*. Bolzano and Weierstrass's definition of what it was for a function to be continuous [*] allows one to *prove* the intermediate value theorem using logic alone. As a matter of logic, any continuous function *in that sense of continuity we have defined*, such that $f(-1) < 0$ and $f(1) > 0$ has some point t between –1 and 1 where it crosses the line — where $f(t) = 0$. The derivation is purely formal, on the basis of definition, without the need for a diagram, picture or any requirement to *see* the conclusion. Advances in mathematics gave the analytic the means to recover lost ground — and to go much further.

It did not go unnoticed that to make sense of the *derivation* of results such as the intermediate value theorem, we needed an expanded notion of logic in order to accommodate reasoning with definitions such as these. The construction "*for every … there is a … such that whenever…*" in the definition of continuity is a complex nesting of what we now understand as *quantifiers*, and it took the development of logic in the 19th and 20th Century in the work of Peano, Frege, Russell and Whitehead to give an account of formal deductive consequence in a vocabulary so expressive. With the expansion of the notion of a definition (to include the Bolzano–Weierstrass definition of continuity) and the expansion of the notion of logic (to include what we now recognise as classical first-order logic), the landscape of the territory between the *a priori* and the *analytic* changed shape.

The development of the tools of logic beyond their Aristotelian bounds gave shape to an important question: where exactly are the bounds of logic? What makes an item of vocabulary a logical constant? This became a live issue with the work of Russell and Whitehead, who appealed to an axiom of *infinity* in the type theory of *Principia Mathematica*. They acknowledged that it was not satisfactory to conceive of this as a logical notion in exactly the same way as other logical axioms. But why? What is the

---

[*] (f is continuous at x if and only if for every $\delta > 0$ there is an $\varepsilon > 0$ such that whenever $|x' - x| < \delta$, $|f(x) - f(x')| < \varepsilon$)

ground for calling something a logical notion? Clearly a claim that there are infinitely many things is not empirically verifiable in any straightforward fashion, so in some sense it may be thought to be *a priori* if it is, in fact, true. But could a logical axiom be the kind of thing we could *debate*? Contemporary discussions of the bounds of logical concepts have not settled on a sharp delineation beyond broad agreement that what is known as 'classical first-order logic' is a uncontroversially within the scope of logic properly so called (see Etchemendy 1990, Read 1994, Shapiro 1991, Sher 1991).

Now, not only did the new mathematics and logic make possible the thought that the truths of calculus could be true 'by definition' without the aid of pure intuition, but advances in geometry and set theory paved the way to do the same thing for space and time, Kant's core notions of geometry and numbers.

Advances in thinking about Euclid's fifth postulate, the *parallel postulate*, led to the construction of formal models of geometric theories in which space behaved radically differently (in the work of Lobachevsky, Reimann and others, and the question naturally arose as to whether these non-Euclidean geometries did any better than Euclid at characterising our pre-theoretic concept of space. Advances in theories of sets, classes and types, in the work of Frege, Russell and Whitehead, and others, meant that such fundamental notions of *number* were having their internal structure plumbed and various analyses of the notions were proposed in order to more clearly articulate the commitments in a theory of number, and to propose various *definitions* of that notion.

At the height of logical positivism, analyses of notions such as *space*, *time* and *number* had been proposed, refined and developed. Carnap's *Aufbau* proposed a logical framework for the definition of concepts, for relating them to sense experience, and for analysing necessity and a priority as due to *analyticity*, which in turn was a purely *conventional* matter. It is up to us which language we use, the adoption of one language over another is an arbitrary or pragmatic matter, and *relative* to that choice of language, a space of necessities and possibilities is defined, which can be analysed *a priori* using the tools of logic. It is up to us that we define logical, numerical, geometrical concepts in the way that we do, in just the same way that we define other terms of our language. Relative to that definition, some truths are necessary and others are contingent, and this matter is purely linguistic, relative to the choice of language employed. Necessity and possibility; *a priority* and *a posteriority* as internal questions are to be answered relative to a linguistic framework. The *external* question, of whether to employ *this* language or *that* one is a pragmatic affair, to be answered not by asking which language better reflects *truth* (for any question of truth is relative to the choice of a language in which that truth is expressed), but rather, by other concerns such as ease of use, better fit with practical or theoretical virtues or other aims of inquiry. The choice, say, between the employment of a non-Euclidean or Euclidean geometry is not the empirical question of which one is *correct*, in the absence of some prior choice of how we are to identify what counts as a *point* and what counts as a *line.* Once we have made that choice — and choice indeed it is, are points abstract, located in experience, are lines the paths of uninterrupted light beams, or to be identified in some other empirical fashion? — then the properties expressed in that language become a matter of empirical or logical investigation, depending on the language chosen at the outset. On this view, the only necessity is *verbal* necessity — necessity grounded in the analytic, the choice of language. What is *a priori* and *necessary* is true by convention, the conventional choice of a linguistic framework.

In this way, by the end of the first third of the 20th Century, the theoretical landscape had inalterably changed. In the analytic tradition at least, Kant's intuition was largely banished in favour of logic, language and convention.

Gödel and Quine

Such an consensus did not survive. The decline of the logical positivists' identification of necessity and a priority with analyticity, and the split between necessary truth-by-convention and contingent truth-on-the-basis-of-reality came in two fronts: one another new mathematical result, and the other, a powerful philosophical metaphor.

The mathematical result is Gödel's justly celebrated incompleteness theorem. Gödel's result dealt a deathblow to the naïve identification of necessary and *a priori* truth (in some language) with what is analytically true (in that language). Gödel showed that in the case of very many mathematical theories — in particular, any theory strong enough to include a small proportion of modern mathematics — there are statements which are true *of* that theory but are not provable *within* that theory, and this can be straightforwardly *proved a priori*. This would not be a problem if it applied to *some* theories, for we could say that even though theory T is incomplete, the theory we are using to reason *about* T is the *real* theory, whose notion of necessity is to be identified with analytic truth relative to *it*. But Gödel's theorem will apply to this theory too, if it is consistent. For any theory (if true), we can construct a stronger theory including truths the first theory missed out. The only mathematically realistic theories to escape Gödel's incompleteness are theories that are *inaccessible* — theories whose axiomatic basis is so complicated as to not be *in principle* enumerated. Such theories are not candidates to be languages in Carnap's sense, for they are not the sort of thing that can be used as frameworks governing our use of mathematical vocabulary. Carnap's program is dealt a deathblow on mathematical grounds.

Not only was Carnap's program dealt a blow by Gödel — these results provide grist for the mill for anyone concerned with the epistemology of mathematics. Which mathematical claims are knowable *a priori*? How are we to account for mathematical truth and our access to that truth? Gödel's results show that this branch of *a priori* truth is a very delicate matter (Franzén 2004, Potter 2000, Smith 2008). One contemporary proposal is to see that numerical vocabulary as introduced by *abstraction* over a relation of equinumerosity. If the Fs are equinumerous with the Gs we say that the number of Fs is the same as the number of Gs (Wright 1983). In this way it could be (in some sense) *a priori*, without reducing to logic, and we do not need to identify the truths of arithmetic with any formal theory, because what we take to be true concerning numbers may well depend on our conceptual apparatus in other matters. (What numbers we are able to countenance depends on what predicates we can construct.) Matters here are, of course, subtle, both philosophically and mathematically (Burgess 2005, Fine 2002).

Gödel's results showed that the *a priori* is not to be identified with the analytic in some language. Quine's arguments against *analyticity* not only led analytic truth into disrepute, but brought down the *a priori* with it. Quine's compelling metaphor of the web of belief, and the attack on the analytic/synthetic distinction apply not only to analyticity but also to the a priori.

In a series of papers, including "Truth by Convention" (1936), "Two Dogmas of Empiricism" (1951), culminating in a book *Word and Object* (1960), Quine led a sustained attack on the logical positivism of his mentor, Carnap. In "Truth by Convention" he argued that the Carnap's distinction between conventional and empirical truth is unsustainable, and along with it, there is no place for the distinction between internal and external questions. In "Two Dogmas," Quine argues that the notion of synonymy required in any account of analytic truth is anything but unrevisable or knowable a priori, and so cannot play the role required of it in the logical positivist programme. Instead of privileging a class of statements as *a priori* and immune from revision on internal grounds (to be changed only on external grounds on the basis of a pragmatic *choice* for one vocabulary over another), Quine argues that the network of beliefs is so interconnected that a difficulty in one area may be fixed by a revision in another, whether that area is close to the periphery of perception statements, some way in at the level of generalisations and lawlike statements, or deep in the centre, at the level of mathematics and logic. All is of a piece, and the entire edifice of commitments stands under judgement from the tribunal of experience *together*. Only as a whole does a theory (an entire epistemic standpoint) serve to be confirmed or disconfirmed by evidence, and any part of it can be revised to better fit that evidence. The epistemology is explicitly *holist*, and *a posteriori*. Putatively *a priori* claims seem independent from experience because they are *general,* applying regardless of the experience we receive from the world and therefore telling us nothing about the world of experience, but nonetheless, perhaps they are revisable in the light of other statements in just the same way as any other claim.

Such a wide-ranging attack on the *a priori* was unprecedented: in Quine's vision of philosophy, none of the notions of *a priority*, analyticity nor necessity play a central role. Philosophy is continuous with the empirical sciences: it differs only in generality.

Kripke and Kaplan

The *a priori*, analytic and the necessary were not without friends in the second half of the 20th Century, despite Quine's attack. Among defenders of the of the notion, closer attention was paid to the relationship between the analytic and the necessary, and this shed new light on possibilities for theories of the *a priori.* Two great insights came from Kripke and from Kaplan.

From Kripke (1972) we learned an example of what may be necessary but not *a priori*. The famous examples are all *identity* statements. Take the claim that Hesperus is Phosphorous, where Hesperus is a name for the morning star, and Phosphorous is a name for the evening star. It is plausible that Hesperus *is* Phosphorous, and that this truth is not merely contingent. That planet (said, pointing to Hesperus) is the very same thing as that planet (pointing to Phosphorous). There is no way that they could differ, for there is no 'they,' it is merely an 'it' pointed to twice. So, says Kripke, we should conclude that the claim that Hesperus is Phosphorous is necessary. But can we know that Hesperus is Phosphorous *a priori*? This seems to not be the case: As a matter of contingent fact, we learned our names in a context in which the Morning and Evening Stars were the same planet. Had things been different, the last body seen in the morning and the first seen in the evening, could have well been different celestial bodies. In that case, the sentence "Hesperus is Phosphorous" would not only fail to be *a priori*, it would fail to be true. The only way we can reassure ourselves that we are in *our* circumstance rather than that circumstance is to to engage in astronomical observations and theorising. We learn that Hesperus is Phosphorous only *a posteriori*.

From Kaplan (1989) we learned the opposite lesson. Not only are some necessities not *a priori,* but some contingent things *can* be known a priori in reverse, not all a priori knowable truths are necessary. I can know, it seems, that I am here now. This claim is true, and in some sense, analytically true. However, it is contingent. Had things been different, I wouldn't have been *here*, I would have been elsewhere. The logic of indexicals, (and demonstratives, such as *this* or *that*) has subtle connections with necessity. The presence of contextual features to be filled in at the circumstances of utterance mean that we can exploit our understanding of these features (or the lack thereof) to generate interesting cases of knowledge, independent of our external evidence.

There are many current issues in these areas: the logic of necessity and names is controversial, and so is the issue of how to extend Kripke's ideas from names to natural kind terms such as *water* and H_2O, as we shall see in the next section. Similarly, when it comes to indexicals and demonstratives, How far are we to go in contextually settled parameters such as that of location, speaker and time of utterance? Does this example mean that we can know statements to be true *a priori*, without knowing what propositions those statements express? What exactly is the item of knowledge claims in these circumstances? What notions of analyticity are in play in these examples? Is the discussion of analyticity merely *ignoring* Quinean objections, or do advances in linguistics and semantics such as those found in Kaplan, Montague and others give us the tools to defuse Quine's objections? (See Russell 2008).

Externalism and Skepticism

Kripke's example has shown us that the behaviour of concepts we have may depend on external factors beyond our immediate grasp, and the ability for us to acquire a concept may depend, in a straightforward manner on external factors. If this is the case, then we seem to have the following puzzle. Consider this argument (Boghossian 1997):

1. If I have the concept water, then water exists.
2. I have the concept water.

Therefore,

3. Water exists.

We seem to know the first premise *a priori*, on general semantic grounds. If I have the concept water, this could only have been acquired in an environment in which someone has been in the presence of water. On the other hand, it is true that we *do* have the concept water. (How could we have formulated this argument if we didn't possess the concept?) It seems to inescapably follow that water exists. But if the premises are known *a priori*, then it seems that the conclusion is known *a priori* too, for it is the conclusion of a valid argument. How can this be? Can we truly know that there is water *a priori*?

This is a stripped down version of an argument raised by Putnam's famous *Brain in a Vat* thought experiment. It seems that we have proved that skepticism is impossible, given the externalist nature of or concepts.

This seems like too much of a success for a priori knowledge: some have thought to resist the argument at the conclusion: to say that warrant is not always transmitted from premises to conclusion of an argument, even when those premises are a priori known (see, e.g., Wright 1991). Others have sought to clarify the premises: while premise 1. is true, if read in the form it appears to take, it cannot be known a priori. What can be known a priori is that

1*. If I have the concept water, then whatever plays the water role exists.

That may not be water, were we to be systematically deceived by a demon, or to be running in a computer simulation. From this premise, all we are entitled to derive is the weaker conclusion 3*, that whatever plays the water role exists, and this is much less surprising as an example of *a priori* knowledge.

A Priori Truth from above & from below

I'll end with a short sketch of two constructive pictures of the nature of *a priori* truth. The first is the framework in which the above disambiguation can be made out: it is the Two-Dimensionalist account of necessity, analyticity and the *a priori*. Originating in work of Humberstone and Davies (1980) analysing the distinction between necessity and analyticity in the presence of operators for *actuality*, the two-dimensional framework has been adopted for a wide-ranging perspective on the connection between necessity — conceived of as truth in all possible worlds (where a possible world is how things could have been *had things gone differently*) and *a priori knowability* as true in all epistemic scenarios (where an epistemic scenario is a way that things as they are *could be taken to be*). In this case, there are epistemic scenarios in which Hesperus is not Phosphorous, and epistemic scenarios in which water is not H_2O, but some other material plays the water role. However, water (the substance that plays the watery role) is of necessity H_2O, so water is H_2O in every possible world.

Understanding what might play the role of an epistemic scenario is not straightforward (Chalmers 2004), but this broad approach to analyticity, necessity and the a priori has motivated a great deal of contemporary philosophy (e.g. Jackson 2000).

This approach to grounding the a priori in truth-in-all-epistemic-alternatives is a 'top-down' *all at once* account of the a priori. It is hostage to giving an account of what is true in an epistemic alternative, and it gives little insight into how that might be found in any concrete case. The main rival to this sort of account of *a priority* is to be found in the *inferentialist* traditions, exemplified currently in the work of Brandom (1994) and Peacocke (2004). In this tradition, we return to the connection between the *a priori* and reasoning: if fundamental inferences involving concepts are to be *a priori*, then perhaps we can use these inferential proprieties to give an account of what it is to possess that concept — indeed, what it is to *be* that concept. To possess the concept of conjunction is, in part, to be disposed to use the inference rules of conjunction elimination (from (p and q) to infer p) and conjunction introduction (from both p and q to infer their conjunction (p and q)). To possess a concept such as the colour concept *green* is to place it in a network of inferences with other colour concepts, as well, perhaps to place it inferentially in a network of *input* and *output* rules governing circumstances where it is permissible to introduce the concept or to exploit it.

Such accounts of the a priori have the great advantage of paying attention to the fine detail of each concept under discussion, and to play an important role in grounding the propriety of inferences employing these concepts. However, they have he great *disadvantage* of requiring a great deal of attention to every single case. Furthermore, it is an open question of whether a non-circular reason can be given for why some concepts may be *introduced* by definition or stipulation and others may rejected as incoherent or defective. If concepts are individuated by their inferential roles, then why can we not justify any questionable inference by the adoption of a concept which just happens to take that inference as one of its defining conditions? (Williamson 2003). Perhaps some combination of a top-down two-dimensional approach and bottom-up inferentialist one will provide resources to respond to these concerns.

As we have seen, since Kant the notion of the *a priori* has waxed and waned in its philosophical fortunes: at some times it has been the centrepiece of philosophical concerns, and at others, it has been peripheral. Just as the relationships between metaphysics, semantics and epistemology vary between philosophical positions and fashions, we may also draw the connections between the notions of necessity, analyticity and the *a priori* (or truth in *all* possible worlds, truth *independent* of experience, and truth determined *only* by meaning) in correspondingly different ways. The different ways we treat the *a priori* reflect broader concerns and larger themes in philosophy, and bring to light our views of the aims and methods of philosophical inquiry.

Reading

There is much to read on current work on the *a priori*. A useful complement to this article is Carrie Jenkins' survey article (2008). It is especially good on the twists and turns of contemporary efforts to define the notion of the *a priori*, and the debate over whether what is *a priori* is unrevisable.

For another, longer general treatment of the current debate, Albert Casullo's article in Moser's *Oxford Handbook of Epistemology* is very clear, as is his 2006 book. For an excellent collection of core readings, Moser's *A Priori Knowledge* is invaluable.

For an entertaining and illuminating account of the *a priori* in Kant, its attempted capture in the work of the logical positivists, and the eventual disintegration of that program, you must read Alberto Coffa's *The Semantic Tradition from Kant to Carnap.*

The best account of Quine's criticisms of analyticity and (implicitly) *a priority* are still Quine himself: both "Two Dogmas" and *Word and Object* are models of clarity. For Gödel's theorems, Smith's new *Introduction* is second-to-none.

The best overview of the mechanics of the two-dimensional account of the *a priori* is to be found in Chalmers' long paper "Epistemic Two-Dimensional Semantics," but a little book that shows how the general approach can be applied to a vast range of philosophical issues is Jacksons *From Metaphysics to Ethics.*

For an introduction to inferentialism, the first port of call must be Brandom's *Articulating Reasons*, but the interested reader must move beyond this short introduction, either to the big book *Making it Explicit*, or to Peacocke's very different *The Realm of Reason.*

Paul Boghossian (1997) "What the Externalist Can Know A Priori". *Proceedings of the Aristotelian Society*, 97(2):161–175.

Robert B. Brandom (1994) *Making It Explicit*. Harvard University Press.

Robert B. Brandom (2000) *Articulating Reasons: an introduction to inferentialism*. Harvard University Press.

John P. Burgess (2005), *Fixing Frege*, Princeton University Press, Princeton.

Albert Casullo (2006) *A Priori Justification*, Oxford University Press, Oxford.

David Chalmers (2004) "Epistemic Two-Dimensional Semantics". *Philosophical Studies*, 118(1):153–226.

J. Alberto Coffa (1993) *The Semantic Tradition from Kant to Carnap*. Cambridge University Press. Edited by Linda Wessels.

Martin Davies and Lloyd Humberstone (1980) "Two notions of necessity". *Philosophical Studies*, 38(1):1–30, 1980.

John Etchemendy (1990) *The Concept of Logical Consequence*, Harvard University Press, Cambridge, Massachussetts.

Hartry Field (2000), "A Priority as an Evaluative Notion," pp. 117–149 in *New Essays on the A Priori*, edited by Paul Boghossian and Christopher Peacocke, Oxford University Press, Oxford.

Kit Fine (2002) *The Limits of Abstraction*, Oxford University Press, Oxford.

Torkel Franzén (2004) *Inexhaustibility*: *a non-exhaustive treatment*, Volume 16 of *Lecture Notes in Logic*, Association for Symbolic Logic.

Robert Hanna (2001) *Kant and the Foundations of Analytic Philosophy*, Oxford University Press, Oxford.

Robert Hanna (2006) *Rationality and Logic*, MIT Press, Cambridge, Massachusetts.

Frank Jackson (2000) *From Metaphysics to Ethics*: *A Defence of Conceptual Analysis*, Oxford University Press, Oxford.

Carrie Jenkins (2008) "A Priori Knowledge: Debates and Developments" *Philosophy Compass* 3(3), 436–450.

David Kaplan (1989) "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals" pp. 481–614 in *Themes from Kaplan*, edited by Joseph Almog, John Perry and Howard Wettstein, Oxford University Press, New York.

Philip Kitcher (2000) "A Priori Knowledge Revisited," pp. 65–91 in *New Essays on the A Priori*, edited by Paul Boghossian and Christopher Peacocke, Oxford University Press, Oxford.

Saul A. Kripke (1972) *Naming and Necessity*. Harvard University Press, Cambridge, Massachussetts.

Paul K. Moser, editor. (1987) *A Priori Knowledge*, Oxford Readings in Philosophy, Oxford University Press, Oxford.

Paul K. Moser, editor. (2002) *Oxford Handbook of Epistemology*, Oxford University Press, Oxford.

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

Michael D. Potter (2000) *Reasons Nearest Kin*: *philosophies of arithmetic from Kant to Carnap*, Oxford University Press, Oxford.

W. V. O. Quine (1936) "Truth by Convention" pages 90–124 in *Philosophical Essays for Alfred North Whitehead*, Longmans, Green and Co., New York. Reprinted in *The Ways of Paradox, and Other Essays.*

W. V. O. Quine (1951) "Two Dogmas of Empiricism". *The Philosophical Review*, 60(1):20–43.

W. V. O. Quine (1960) *Word and Object*. MIT Press.

Georges Rey. "A Naturalistic A Priori". *Philosophical Studies*, 92(1):25–43, 1998.

Gillian Russell (2008) *Truth in Virtue of Meaning*, Oxford University Press, Oxford.

Stewart Shapiro (1991), *Foundations without Foundationalism*: *a case for second-order logic*, Oxford University Press, Oxford.

Gila Sher (1991), *The Bounds of Logic*, MIT Press, Massachussetts.

Peter Smith (2008), *An Introduction to Gödel's Theorems*, Cambridge University Press, Cambridge.

Timothy Williamson (2003) "Blind Reasoning". *Aristotelian Society Supplementary Volume*, 77(1):249–293.

Crispin Wright (1991) "Scepticism and Dreaming: Imploding the Demon." *Mind* 100(1):87–116.