

Williamsonian Scepticism about the A Priori

Giacomo Melis and Crispin Wright

We focus on Timothy Williamson's recent attack on the epistemological significance of the a priori–a posteriori distinction,¹ and offer an explanation of why, fundamentally, it does not succeed. We begin (§1, 2) by setting out Williamson's core argument, and some of the background to it and move to consider (§3, 4) two lines of conciliatory response to it—conciliatory in that neither questions the central analogy on which Williamson's argument depends. We claim, setting aside a methodological challenge (§5) to which Williamson owes an answer, that no satisfactory such reconciliation is in prospect. Rather, as we then argue, —and it is on this that we base our overall negative assessment of his argument—Williamson's core analogy is flawed by an oversight (§§6-8). We conclude (§9) with some brief reflections on Williamson's ideas about the imagination as a source of knowledge.

Our principal conclusion is only that Williamson's argument fails to perform as advertised. A constructive case for confidence, to the contrary, that the intuitive contrast between a priori and a posteriori reflects something of fundamental epistemological significance is prefigured in our final section, but will not be elaborated here.

§1 Background

The history of philosophy attests that there has seemed to many to be a genuine distinction between judgements warranted by empirical evidence and those warranted 'a priori'. Scepticism about whether anything substantive is marked by this tendency has often — maybe usually—been driven by background theoretical assumptions. The holistic empiricism of Quine's "Two Dogmas" is one salient example. Another is epistemological externalism. John Hawthorne, for example, has written that

If an epistemological distinction fails to carve at the epistemological joints, then it is not worthy of serious and protracted discussion. The residual issue whether the putative distinction is incoherent or merely gerrymandered ought not to strike anyone as especially important. My own externalist commitments – epistemological and semantic – lead me to think that the a priori / a posteriori distinction is not a practically natural one, and hence that its importance to epistemology has been grossly overestimated.²

Hawthorne goes on to argue that, from an externalist standpoint,³ there are serious difficulties with each of two kinds of common characterisation⁴ of a priori knowledge in capturing the intuitively desired extension. However, even his argument is successful in

¹ Williamson (2013). All numbered page references are to this text unless otherwise stated.

² Hawthorne (2007: 201)

³ Specifically, from a standpoint that regards *safety*—where, roughly, a true belief is formed safely just if formed by the method actually used, it could not easily have been mistaken—as necessary for knowledge.

⁴ Respectively, *environment-independence* —the idea that a process warrants a belief a priori just if its execution would provide a warrant for that belief no matter what the environment—and *experience-independence*—the idea that warrant for a belief is a priori provided it is acquired in a way that does not essentially involve sense-experience.

its own terms — something which we offer no opinion about here—it merits emphasis that any externalistically-driven line of attack on the a priori is tendentious. The classical a priori, encompassing judgments about valid inference, basic arithmetic, geometry, modal status, the phenomenology of colour, and so on, may turn out to be a gerrymandered category, but the common phenomenology of cognition in all the basic examples is one of reflective intellection, perhaps aided by thought-experiment and imagination. Persuading oneself that seven plus five is twelve, that no surface can be simultaneously red and green all over, or that modus ponens is valid, seems to be something we can do in a perfectly epistemically responsible fashion in the armchair, with our eyes closed, purely by the exercise of certain intellectual and imaginative routines. In short, the classical a priori supplies perhaps the most natural home for the fostering of *internalist* conceptions of epistemic warrant—conceptions that hold that the attainment of warrant is grounded in internal mental processing of which agents can be fully aware in normal self-consciousness and for which they can take full responsibility. Such conceptions may indeed prove to be problematic when generalised. But the dialectical position is that wherever internal conscious mental processing is ordinarily recognised as an *unconditional*—though maybe defeasible—source of epistemic warrant, then that is a strike in favour of at least a local epistemic internalism. The credibility of a general epistemic externalism thus depends in its being able either to account for the a priori, or to discredit it independently. Doubts about it that presuppose externalism—which in effect merely draw on problems associated with making sense of it in externalist terms—ought to be regarded as dialectically ineffective.

Williamson, like Hawthorne, is of course of well-documented, strong externalist sympathies. But one consideration that makes his critique of the a priori/a posteriori distinction potentially much more interesting is that he does not present his central argument as relying on any externalist assumptions but at least appears to aim at a broad theoretical neutrality.⁵

§2 *Williamson's critique outlined*

There is another consideration that adds to the interest of Williamson's argument: the fact that he is not sceptical about the *reality* of the a priori/a posteriori distinction. He is quite content to grant that there is a distinction, which people are capable of picking up from examples and applying non-collusively to new cases in ways which produce broad agreement with each other – in short, that there is, as he puts it, a “bottom-up” route into

⁵ We say “appears” because there is no passage in the paper that explicitly expresses this aim. But the following footnote (number 2, p. 310), though it flags a switch of attention from knowledge to justification — as if attention to the latter was somehow the province of internalism—may be thought indicative:

Many contemporary epistemologists, especially those of an internalist bent, treat the [a priori/a posteriori] distinction as primarily a classification of forms of justification rather than of knowledge. For reasons explained in Williamson (2000), I treat the classification of forms of knowledge as primary. Friends of justification should not find much difficulty in reworking the arguments of this chapter in their terms.

We will return to the question whether Williamson's' dialectic really is free of externalist presupposition later.

capacities of productive, non-collusive, broadly consensual judgements about what is a priori and what is a posteriori.

The thrust of Williamson's critique is twofold: first, that the profile of the distinction which we thus draw bottom-up does not coincide with any theoretically important or explanatory epistemological distinction; second, that standard official characterisations of the distinction — characterisations that proceed “top-down”, as he likes to say — misrepresent the role played by experience in the kind of cognitive processing characteristic of (many?) a priori judgements, and fail to align with the distinction as drawn bottom-up. Like Hawthorne, Williamson characterises the general thrust of his critique as to the effect that the traditional a priori/a posteriori distinction fails to “carve at the epistemological joints”.

To elaborate. The traditional understanding of the a priori/a posteriori distinction proceeds on the assumption, commonly accepted, that experience can play two quite different roles in knowledge and justification: one *evidential*, the other *enabling*. Experience plays an evidential role when it provides the agent with her evidence for a certain judgement; it plays a merely enabling role when it is required only for the acquisition of the concepts drawn on in an understanding of the proposition that she is otherwise warranted in judging true. Since the justification of any proposition draws on an understanding of the concepts configured in it, and since concept-acquisition is an empirical process, defenders of the a priori are therefore normally willing to grant that justification a priori is not justification which is “independent of experience” in its enabling role. Their thought is rather that, in such justification, experience plays no *evidential* role.

So much is familiar. Williamson considers the following two propositions:

- | | |
|--------------------|--|
| <i>(Crimson)</i> | All crimson things are red |
| <i>(Who's Who)</i> | All recent volumes of <i>Who's Who</i> are red |

The traditional view is that the role played by experience in acquiring justification for *(Crimson)* is merely enabling: experience is required to form the concepts, *crimson* and *red*. But once equipped with the concepts, so runs the traditional thought, agents are able to acquire justification for *(Crimson)* simply by reflection on the colours involved, without any need for additional experience, in particular, without having to seek out crimson things and verify whether they are red or not. By contrast, it doesn't seem that an agent can get justification for *(Who's Who)* without checking the colour of recent volumes of *Who's Who*. In addition to the experiences required for the formation of the concepts involved in *(Who's Who)*, some kind of empirical investigation is called for. The heart of Williamson's argument is a direct challenge to this way of looking at these examples.

Williamson imagines a character — Norman — who is trained ostensibly in the use of the words, "red" and "crimson", independently of each other, encountering both positive and negative samples of each. We are to suppose that his training goes well and that he becomes expert in the simple classificatory use of both expressions. His training does not involve any consideration of the proposition, *(Crimson)*, and now he is invited to take a view about it. In Williamson's view, he can competently do so purely by an exercise of the capacities of visual imagination, which, it may plausibly be supposed, his

training will have fitted him with: he has only visually to conjure up a sample of crimson and then “within the imaginative supposition” recognise that it is indeed red. No recollection of any specific previous encounters with red or crimson things is needed. And since Norman’s performance may be supposed to have been “sufficiently skillful”, and the background conditions to it to have been perfectly normal, his consequent judgement that (*Crimson*) is true should count as knowledgeable.

This may seem to pass muster as an outline of how (*Crimson*) and similar propositions may be recognised to be true a priori in something like the 'bottom-up' understanding of the notion.⁶ The sting comes when, according to Williamson, we note that exactly the same kind of reflective, imaginative routine is, in relevantly similar circumstances, at the service of knowledge of a 'bottom-up' a posteriori proposition like (*Who's Who*). We have only to suppose that Norman has successfully undergone a similar ostensive teaching in the correct classificatory use of the expressions, “red” and “recent volume of *Who's Who*” — similar in that he has encountered both positive and negative examples of the application of each, that they have been explained independently of each other, and that the teaching has involved no explicit consideration of (*Who's Who*). In order to come to a judgement of that proposition, Williamson avers, all Norman has to do is, similarly, to conjure up in his visual imagination a sample of a recent volume of *Who's Who* and correctly judge the colour that it is therein represented as having. Once again, provided the performance is sufficiently skilful, and assuming that the background conditions are normal, there is no reason to deny that Norman's judgement is knowledgeable.⁷

⁶ Below (§6), we will question whether it really does pass muster.

⁷ For the benefit of readers without immediate access to Williamson’s text, it may be helpful to quote the relevant passage in full:

Suppose that Norman acquires the words 'crimson' and 'red' independently of each other, by ostensive means. He learns 'crimson' by being shown examples to which it applies and samples to which it does not apply, and told which are which. He learns 'red' in a parallel but causally independent way. He is not taught any rule like (*Crimson*), connecting 'crimson' and 'red'. Through practice and feedback, he becomes very skilful in judging by eye whether something is crimson, and whether something is red. Now Norman is asked whether (*Crimson*) holds. He has not previously considered any such question. Nevertheless, he can quite easily come to know (*Crimson*), without looking at any crimson things to check whether they are red, or even remembering any crimson things to check whether they are red, or making any other new exercise of perception or memory of particular coloured things. Rather, he assents to (*Crimson*) after brief reflection on the colours crimson and red, along something like the following lines. First, Norman uses his skill in making visual judgments with 'crimson' to visually imagine a sample of crimson. Then he uses his skill in making visual judgments with 'red' to judge, within the imaginative supposition, “It is red”. This involves a general human capacity to transpose ‘online’ cognitive skills originally developed in perception into corresponding “offline” cognitive skills subsequently applied in imagination. That capacity is essential to much of our thinking, for instance when we reflectively assess conditionals in making contingency plans (See Williamson 2007: 137-78). No episodic memories of prior experiences, for example of crimson things, play any role. As a result of the process Norman accepts (*Crimson*). Since his performance was sufficiently skilful, background conditions were normal, and so on, he thereby comes to know (*Crimson*).

Naturally, that broad-brush description neglects many issues. For instance, what prevents Norman from imagining a peripheral shade of crimson? If one shade of crimson is red, it does not follow that all are. The relevant cognitive skills must be taken to include sensitivity to such matters. If normal speakers associate colour terms with central prototypes, as many psychologists believe, their use in the imaginative exercise may enhance its reliability. The proximity in colour space of prototypical crimson to prototypical red is one indicator, but does not suffice by itself, since it does not discriminate between “All crimson things are red” (true) and “All red things are crimson” (false). Various cognitive mechanisms can be postulated to do the job. We need not fill in the details, since for present purposes what matters is the overall picture. So far, we may accept it as a sketch of the cognitive processes underlying Norman's a priori knowledge of (*Crimson*).

So, Williamson is offering an example where, so he claims, the very same type of cognitive process provides justification, and indeed knowledge, for both (*Crimson*), that all crimson things are red, and (*Who's Who*), that all recent volumes of *Who's Who* are red. If he is right, he has shown, at the least, that there is no important cognitive difference between one way of knowing the usually regarded a priori proposition, (*Crimson*), and one way of knowing the usually regarded a posteriori proposition, (*Who's Who*). But of course there is an asymmetry in the traditional understanding of the notions. While an a posteriori proposition, traditionally regarded, is one which can *only* be known by, broadly speaking, empirical means, an a priori proposition, while it can be known a priori, may also be knowable a posteriori. So before Williamson's example, even if we find no problem with his description of it, can put any kind of pressure on the distinction, it needs to be argued that Norman's cognitive processes as described are not properly regarded as constituting an *a posteriori* mode of knowledge-acquisition. We'll come back to this.

Williamson wants the example to show two things. Remember that he accepts that there is an intelligible “bottom-up” distinction between the a priori and the a posteriori. Accepting such a distinction, the first intended point of the example is that there need be very little difference in what goes on when, respectively, an agent competently comes to know, a priori, that all crimson things are red and when he comes to know the, by bottom-up standards, a posteriori proposition that all recent copies of *Who's Who* are red, and hence very little difference between (one kind of) putatively a priori cognitive processing and (one kind of) a posteriori cognitive processing. This point, again, depends on accepting that what Norman, as described, does in coming to know (*Crimson*) does properly count as coming to know it a priori.

The second point Williamson intends the example to make is that the traditional distinction between the evidential and the enabling role of experience in cognition misses an important *third way* in which experience can be involved in the acquisition of justification. This third way is tied to the knowledge-productive use of imagination in the

Now compare the case of (*Who's Who*). Norman is as already described. He learns the complex phrase “recent volumes of *Who's Who*” by learning “recent”, “volume”, “*Who's Who*” and so on. He is not taught any rule like (*Who's Who*), connecting “recent volume of *Who's Who*” and “red”. Through practice and feedback he becomes very skilful in judging by eye whether something is a recent volume of *Who's Who* (by reading the title), and whether something is red. Now Norman is asked whether (*Who's Who*) holds. He has not previously considered any such question. Nevertheless, he can quite easily come to know (*Who's Who*), without looking at any recent volumes of *Who's Who* to check whether they are red, or even remembering any recent volumes of *Who's Who* to check whether they are red, or any other new exercise of perception or memory. Rather he assents to (*Who's Who*) after brief reflection along something like the following lines. First Norman uses his skill in making visual judgments with “recent volume of *Who's Who*” to visually imagine a recent volume of *Who's Who*. Then he uses his skill in making visual judgments with “red” to judge, within the imaginative supposition, “It is red”. This involves the same general human capacity as before to transpose “online” cognitive skills originally developed in perception into corresponding “offline” cognitive skills subsequently applied in imagination. No episodic memories of prior experiences, for example of recent volumes of *Who's Who*, play any role. As a result of the process Norman accepts (*Who's Who*). Since his performance was sufficiently skilful, background conditions were normal, and so on, he thereby comes to know (*Who's Who*).

As before, the broad-brush description neglects many issues. For instance, what prevents Norman from imagining an untypical recent volume of *Who's Who*? If one recent volume of *Who's Who* is red, it does not follow that all are. The relevant cognitive skills must be taken to include sensitivity to such matters. As before, Norman must use his visual recognitional capacities offline in ways that respect untypical as well as typical cases. We may accept that as a sketch of cognitive processes underlying Norman's a posteriori knowledge of (*Who's Who*). Williamson (2013: 295-296).

kind of "offline" application of skills acquired through perception that Norman is meant to illustrate. Continue to allow that a way has indeed been described whereby Norman can come to know that all crimson things are red. Williamson contrasts Norman's case with that of Norbert, a competent native speaker of English who has acquired the words "red" and "crimson" in a way similar to Norman, but has not yet had very much practice and feedback in classifying things as "crimson" or "not crimson". We can imagine, though, says Williamson, that Norbert usually gives the correct verdicts with both "red" and "crimson" at least when applying his skill of making colour judgments "online"—that is, in response to real external cases. So, the suggestion is, we can suppose him linguistically competent with both words: he has the concepts, *red* and *crimson*, and there is thus no question but that he grasps the proposition, (*Crimson*). However we are to suppose that Norbert is relatively inexperienced in the application of these concepts, so that though competent, he is less skilful than Norman when it comes to *imagining* a crimson sample; and thus he is not able to convince himself that (*Crimson*) is true just by an exercise of the visual imagination. We can suppose that he comes to no view about the issue. But now, since both Norbert and Norman have had the experience necessary to grasp the relevant concepts but only Norman has managed to get to know (*Crimson*) by imagining a crimson sample, it seems that we have to say that the role of experience in Norman's acquisition of knowledge is more than the merely enabling role that it also plays for Norbert. The merely enabling role of experience is not enough to underwrite the acquisition of knowledge of (*Crimson*) by means of the visual imagination. Something more is needed, something that Norman has but Norbert lacks. What Norbert lacks and Norman has is *sufficient* experience to "hone and calibrate his skills in applying the terms "crimson" and "red" to the point where he could carry out the imaginative exercise successfully".⁸

So the role of experience in Norman's route to knowledge of (*Crimson*) is more than enabling, as that is traditionally understood. But that it is also less than evidential simply follows from the description of the case: any direct contribution by perception, memory, testimony or any other experiential source was excluded from the start. Given the close analogy between the cognitive processing involved in (*Who's Who*) and that involved in (*Crimson*), Williamson thus takes himself to have shown that the role of experience in both cases is more than purely enabling and less than strictly evidential.

Williamson's contention, accordingly, is that the traditional distinction between cognitive processes in which experience plays an evidential role and cognitive processes in which it plays a merely enabling role is too crude — that there are significant cognitive processes, culminating in knowledge, in which the role of experience is more than enabling but less than evidential. Hence the traditional kind of "top-down" explanation of the distinction between the a priori and the a posteriori is flawed. Moreover some cognitive processing culminating in what, by bottom-up standards, counts as a priori knowledge is effectively identical to the kind of processing involved in coming to what, by bottom-up standards, count as examples of a posteriori knowledge. The traditional distinction between the a priori and the a posteriori, explicated as it standardly is, thus "fails to carve at the epistemological joints".

⁸ Williamson (2013: 297).

§3 Conciliatory Responses to Williamson's Critique (I): the Synthetic A Priori

What should we think about this? A first reaction is that Williamson needs to show that the way he outlines of knowing (*Crimson*) is in some way exemplary – that it represents, as it were, not merely a legitimate but a *prototypical* way whereby knowledge a priori, both of (*Crimson*) and other central examples of the traditional a priori, can be achieved. If there are other ways of knowing (*Crimson*) which are also plausibly a priori and are not tantamount to Norman's way, then it seems that the theoretical utility of the traditional distinction need be under no immediate challenge by the example, and might yet be vindicated. We could allow that Williamson has shown that the role which experience can play in knowledge-acquisition is more complex than allowed for by the enabling/evidential dichotomy and that we need to recognise the, as it were, intermediate role that experience can play in honing the skills involved in “offline” exercises of the imagination. But we could propose in response that Williamson's observation calls for complication, not rejection; that we need a *tripartite* division. There is the traditional a priori: knowledge acquired by means in which experience plays only an enabling role. There is the traditional a posteriori, in which experience plays both an enabling and an evidential role. And there is a third category in which experience plays both an enabling and the third, intermediate role described by Williamson, but plays no evidential role.

It's in keeping with the spirit of this response to grant that the visual imagination is indeed a cognitive skill that draws on experience but does not deploy the evidence of experience in any straightforward way. And this, for all we are going to argue to the contrary, may be an important insight. The idea that experience can “hone” the use of the conceptual skills we employ in the imagination, and in consequence can put agents in a position to conduct and learn from thought experiments in various ways that would not otherwise be feasible, should seem well worth exploring. One interesting issue in this connection might be whether there are propositions, or subject-matters, that are characteristically associated with knowledge achieved by the intermediate route, in the way in which the evidence of experience is characteristically associated with ordinary empirical propositions, and the deployment simply of conceptual grasp is associated with the traditional analytic propositions. Perhaps Williamson has reminded us, in short, of something characteristic about the methods of cognition that are relevant to the traditional *synthetic a priori*. In that case, it will simply have been a mistake — traceable perhaps to the legacy of positivism and the empiricists — to think of the a priori as exhausted by methods of cognition that demand only conceptual reflection or acquaintance with meanings.

This accommodating response, then, would hold that the effect of Williamson's argument, if sustained, should be to remind us that the traditional way of thinking about the a priori that comes down to us from the positivists, and more generally, the “linguistic theory” of the matter popular in the mid-twentieth century, is seriously impoverished: that a richer epistemology needs to be invoked to accommodate the kinds of cognitive routine – imagination, thought-experimentation, and whatever else careful attention the phenomenology of basic a priori conviction might disclose – involved in the full range of instances of the a priori that we allow “bottom-up”. Propositions like “ $3+2=5$ ”, when suggested by a diagram, or “All cubes have twelve edges”, when suggested by a count of the edges of an imaginary cube, or “Nothing can be both red and green all over”, when

suggested by the phenomenology of colour experience, all come within that range. The picture that would emerge would then be one in which the traditional domain of the a priori would in effect be split in two: on the one hand come propositions like “All bachelors are unmarried”, typically knowable simply by analytic reflection, and in which the use of the visual or other forms of imagination plays no essential role; and on the other hand come propositions like “All cubes have twelve edges” which may be knowable by pure conceptual reflection and deduction from definitions, but which also, and indeed perhaps more usually, are known with the help of visualisation and attention to the figures visualised.

A more satisfactory taxonomy, it may be suggested, might accordingly involve three types of propositions:

- (i) Propositions that can be known only through processes in which experience plays an evidential role – ordinary empirical propositions;
- (ii) Propositions that can be known through methods in which experience is involved merely in an enabling, concept-supplying role; — the traditional analytic a priori; and
- (iii) Propositions that cannot be known merely on the basis of grasp of the concepts involved but can be known without reliance on experiential evidence, by routines which involve essential play with thought-experimentation or imagination and which rest on experience only in the intermediate role that Williamson gestures at.

Of course this response to Williamson’s argument is very much not in the spirit of his own response. Still, it might be prompted by his reflections and it might represent a correct step in epistemology even if his argument is wrong.⁹

However it is, as the reader will likely be impatient to point out, untenable if he is right. The problem with it is, of course, its failure to take the measure of the putative applicability of the method of cognition described by Williamson to the proposition (*Who's Who*), that all recent volumes of *Who's Who* are red. It is surely just too much of a stretch of the intent of the traditional distinction to allow that such a matter can be known a priori, that it is an example of the synthetic a priori. It's one thing to grant that the domain of the analytic incorporates too narrow an epistemology of the a priori, and overlooks the intermediate role that experience-honed imagination can play in furnishing genuine reflective knowledge. To grant that much does indeed call for no more than complication of the traditional distinction. But the distinction is surely not merely complicated but *subverted* if we wind up counting Norman's knowledge of (*Who's Who*) as a genuine species of a priori knowledge as well. Indeed, although Williamson's official line is that his argument comes not to bury the traditional distinction but merely to question its theoretical significance, he surely understates its force. If there really is no relevant epistemological difference between Norman's ways of coming to know (*Crimson*) and (*Who's Who*) as described, the conclusion has to be that in a wide class of cases, the distinction between the a priori and the a posteriori is a distinction without a difference, that it is no distinction at all. From this perspective, Williamson’s concessiveness to the bottom-up distinction — his willingness to allow that there is a distinction, although one of no theoretical or explanatory interest, and permitting of no

⁹ See Jenkins/Kasaki (2015: §§ 4-5) for some suggestions about an analogy between the roles that Kant and Williamson assign to imagination, and the possible complication of the enabling/evidential distinction.

systematic account — may impress as a relatively (uncharacteristically!) modest response to his argument. Everything else being sound, the conclusion seems justified that there is no real distinction: that, in our bottom-up judgments about the a priori and the a posteriori, while we largely march in step with each other, we track no objective divide.

§4 *Conciliatory Responses to Williamson's Critique (II): Extending the A Posteriori*

So far we have been concerned with responses to Williamson's argument which grant that there is a good analogy between Norman's respective ways of knowing (*Crimson*) and (*Who's Who*), but incline to regard that way of knowing, generally characterised, as belonging in spirit with the a priori as traditionally conceived. There is, of course, as briefly signalled earlier, the opposite way of receiving Williamson's case: allow the analogy but propose, instead, that what it shows is that some apparently reflective methods of cognition are actually better classified as *a posteriori* — that knowledge that all crimson things are red, when achieved in Norman's way, is, no less than knowledge achieved in that way of (*Who's Who*), a kind of a posteriori knowledge.

One way of elaborating this response¹⁰ is to argue that Norman's successfully conjuring an image of a crimson shade, or of a volume of *Who's Who*, should be reckoned to be an instance of *recollection* — that is, as something involving memory of what crimson things, and volumes of *Who's Who*, look like. (There would, to be sure, still be a difference in the modal status of what was thereby recollected — crimson things necessarily look like that; volumes of *Who's Who* do not. But almost everyone now allows that necessary propositions can be known a posteriori.) It is true that no *episodic* memory need be involved — no memory of any particular occasion on which a crimson thing, or a volume of *Who's Who*, was encountered—but there certainly is such a thing as remembering, based on acquaintance, what something, or some kind of thing, looks like. And in general such memory need involve no recollection of any particular episode of acquaintance with the thing or with an instance of the kind. If, in a temporary total amnesia, you lose all remembrance of the events of your past but still find yourself able to find your way to your house and recognise it as such, that is knowledge based on acquaintance and stored in memory. Isn't it knowledge of this genre that plausibly underpins Norman's achievements? It is thus questionable whether Williamson produces any real argument for the claim that memory is not involved in Norman's successes. And if it is allowed to be so, then — viewing memory simply as a retrospective recovery of information once yielded by experience — the claim that Norman's method of cognition is a posteriori even in the case of (*Crimson*) may seem to have legs.¹¹

¹⁰ A different way of developing a response of the same shape is proposed by Albert Casullo (this volume). We think that Casullo's proposal also succumbs to the objection to be lodged below.

¹¹ Expectably, Williamson anticipates an appeal to generic memory. He writes with respect to (*Crimson*): "One interpretation of the example is that, although Norman's knowledge of (*Crimson*) does not depend on episodic memory, and he may even lack all episodic memory of any relevant particular colour experiences, he nevertheless retains from such experiences generic factual memories of what crimson things look like and of what red things look like, on which knowledge of (*Crimson*) depends". But then he counters that "such reclassification is a risky strategy for defenders of the a priori-a posteriori distinction. Instead, it might be proposed [...] that the only residue of Norman's colour experience active in his knowledge of (*Crimson*) may be his skill in recognising and imagining colours." Likewise with respect to (*Who's Who*), he writes: "Norman's knowledge of (*Who's Who*) can be envisaged in parallel to his knowledge of (*Crimson*) as just envisaged. [...] The only residue of his experience of recent volumes of *Who's Who* active

The distinction is well taken, but does it help? The real problem with any “Both are a posteriori” response to Williamson’s analogy is the shrinking that it threatens of the domain of the traditional a priori. In particular, if appeal to the reflective imagination is to count as conferring a priori knowledge only in cases where no episode-*unspecific* recollection is involved of what items are like which instantiate the concepts that the imagination makes play with, how many of basic arithmetical, geometrical, and a host of other types of claim are going to survive as knowable a priori? Even the apriority of traditionally acknowledged analytic claims will come under threat in any cases where understanding their key ingredient concepts involves knowing— generic memory— what things that fall under them “are like”. Perhaps we might yet be able to stabilise some notion of the a priori within the boundaries grudgingly recognised by J.S Mill – the trivially verbal and definitional. And those who, contra Williamson, incline to think of the a priori as an important epistemological category might still be right after such a restriction —for recognition of truth based purely on the grasp of meanings does seem to be a real phenomenon, and is not easily explained. But proponents of the classical a priori surely didn’t intend the notion to have so meagre an extent.

§5 *A Methodological Aside*

At this point, though, there may seem to be something methodologically suspect about the argument. Let's take stock. The classical account was that a priori propositions are those that may be known in a way which depends on experience only in a concept-enabling role. This failed to get the extension of the bottom-up distinction right: Williamson found cases where the role of experience is not merely enabling, but which the bottom-up distinction classifies as a priori. The accommodating response then said: Count these Williamsonian cases as *synthetic* a priori, and broaden the category of the a priori accordingly. In response, Williamson can point out that in that case “All recent copies of *Who’s Who* are red”, as known by Norman, counts as (synthetic) a priori too. That seems too much to accept. So the accommodating response, in its turn, misdescribes the extension of the bottom-up distinction. And now we are invited to conclude that the bottom-up distinction is no objective distinction, that it admits of no systematic theoretical account, and is fit for no explanatory work. But why should we draw this conclusion? It is certainly not true in general that when philosophical accounts are proposed to rationalise and explain a distinction that we draw intuitively, and then prove repeatedly to over- or under-extend, we conclude that the original distinction is empty, or of no real theoretical interest. The more common— and certainly another possible— reaction is to conclude that the trouble lies at the level of the proposed accounts. That indeed is exactly the conclusion that Williamson himself has drawn in response to the long line of counter-examples to various proposals about knowledge spawned by the original Gettier cases. There his own recommended conclusion is that there is in the end no saying, at a reductive level, what knowledge consists in: that it should be taken as an epistemological primitive. That is a far cry from Williamson's conclusion in the present instance. Williamson concludes that the a priori is theoretically uninteresting; whereas in the case of knowledge his conclusion is that the notion of knowledge is primitive and

in his knowledge of (*Who’s Who*) is his skill in recognising and imagining such volumes.” (All quotes from Williamson 2013 at p. 298). We agree that the strategy is “risky” — for the reason we are about to explain.

fundamental in epistemology. What breaks the tie? Why doesn't Williamson simply conclude that we presently have no adequate account of the a priori, and that we need to try harder — but mindful that we may have in the end to conclude that this notion too gestures at something primitive and fundamental in epistemology?

It's a good question. When all attempts at a theoretical reconstruction of a notion seemingly fail to get the intuitive extension right, there is always the option of concluding that there is underlying our use of it no consistent, proper rationale — that any broad consensus exhibited in our practice with it is to be explained not in terms of our detection of some real property of (epistemic) situations, but in other terms. (Deflationary and expressivist accounts offer just such explanations.) Granted, if one took this view of knowledge, there would still be a "bottom-up" notion, manifest in our practice with the term. But Williamson must allow that, if no more is said, the response to the post-Gettier history is available that *knowledge* is without explanatory interest, something of which no account is to be given, and that serious epistemological work should employ other notions. This is exactly the conclusion that he wants to draw about the a priori. What has the a priori done, in Williamson's view, to deserve such a conclusion, when the notion of knowledge, in his view, doesn't?

§6 *Questioning the Core Analogy (I)—worries about (Crimson).*

So far, we have been concerned with the significance of, and possible responses to, Williamson's analogy on the assumption that it should be sustained: that imaginative reflection, unaugmented by additional empirical evidence, can be equally productive of knowledge of (*Who's Who*) as of (*Crimson*). Our finding so far is that neither of two possible conciliatory responses seems satisfactory. But should Williamson's analogy be sustained in the first place? It is important, obviously, that the parallel between Norman's cognition in the two cases be given in sufficient detail to make it plausible not merely that he has done much the same thing but that in both cases he has done something by which he is properly convinced, and which — if nothing else goes wrong — could properly constitute the acquisition of knowledge.

It is here that, on closer inspection, the shoe begins to pinch — that the analogy begins to seem procrustean. To begin with, Williamson's description of the case of (*Crimson*) leaves it completely unclear what role the imagination plays in warranting the step from an appreciation of the impression of the visualised patch of crimson to the generality of the judgement concerning all crimson things. As a first step, one could say that, before any belief about crimson things in general can be warranted by this process, one needs some kind of assurance that the patch of crimson visualised is properly *stereotypical* of crimson — that it somehow typifies all shades of crimson. And this information does not seem to be something that merely imagining a particular shade can deliver: rather, Norman will need some independent grip on the *spread* of crimson in the colour band, as it were, before he can be in position to be assured that the shade he visualises is in some way central or typical of them.

Williamson of course acknowledges a worry in this vicinity.¹² His emphasis at this point on Norman's "appropriate cognitive skills" suggests he might reply by insisting

¹² "Naturally, that broad-brush description neglects many issues. For instance, what prevents Norman from imagining a peripheral shade of crimson? If one shade of crimson is red, it does not follow that all are. The relevant cognitive skills must be taken to include sensitivity to such matters. If normal speakers associate

that there is no evident requirement that Norman should *know* that the visualised patch is stereotypical – that it is enough merely that it be so, and that this can be ensured not by information that he possesses but by his well-honed imaginative skills. *If*, though, that is the response that Williamson himself would favour, then we — and he — should remember that we are here being presented with an argument that is supposed to be neutral on the opposition between externalist and internalist conceptions of knowledge and justification. Such a response is not neutral: it depends on a reliance upon the idea that the exercise of a certain kind of reliable imaginative skill can be sufficient to underwrite knowledge even in circumstances where the agents themselves have no reason to suppose that they do indeed have such a skill.

Setting that aside, however, there is further cause for dissatisfaction with Williamson's description of the case. It is crucial to his purposes that Norman fully grasps the concepts of *crimson* and *red*, and yet prior to the episode in his imagination, that it have been an open question to him whether all crimson things are red. The question is how, merely by imagining what is as a matter of fact a stereotypical shade of crimson, he somehow succeeds in closing the question off. What is it about the process that mandates his conviction? Why, now that for the first time he considers the question, and having visualised what is as a matter of fact a stereotypical shade of crimson, is Norman rightly inclined to discount the possibility that there might be other shades of crimson that are not red? How does the imaginative exercise make it plausible to him that there are no such shades?

Williamson himself said, in setting up the case, that it has been *by practice and feedback* that Norman has become skilful in judging by eye whether something is crimson and whether something is red. What kind of feedback did he have in mind? Does it include cases where Norman has overextended the use of “crimson” among red things and been corrected? Williamson himself seems to imply so.¹³ But if *that* kind of thing has been Norman's experience, then his coming to knowledge that all crimson things are red doesn't need the imaginative exercise described; rather it can be grounded in his linguistic experience – he has been advised, by example and correction, that correct use of “crimson” is restricted to a band, vaguely demarcated no doubt, within the determinable colour red. Maybe the episode in imagination plays a causal, or triggering role – maybe it reminds Norman in some way of the constraints on proper practice with “crimson”; but it is hard to see that it stands to teach him anything.

We had better suppose, therefore, that Norman has had no such feedback. He has merely become adept, by practice, in applying “crimson” to crimson things and “red” to red things. But how in that case, merely by visualising a typical case of crimson and, as it were, observing its redness, can Norman reach a rational conviction that all crimson things are red? Something epistemologically essential is clearly missing here: suppose the question had been, conversely, whether all red things are crimson, and that Norman had done exactly the same in response – imagined a shade of red that, as it might well be, was

colour terms with central prototypes, as many psychologists believe, their use in the imaginative exercise may enhance its reliability.” (Williamson 2013: 296), quoted in n. 7 above.

¹³ He writes of Norman's education (p. 295), recall, that “he learns ‘crimson’ by being shown examples to which it applies and samples to which it does not apply, and told which are which.” It is not explicit, though, whether the negative samples for “crimson” included some red ones.

crimson; crimson is a pretty typical shade of red after all. What stops him making the incorrect, converse generalisation? That would not be a well-founded conviction. But Williamson has said nothing, in his account of Norman's cognitive processing, to account for the difference.

Again, Williamson seems fully aware of exactly this. Recall what he says:

The proximity in colour space of prototypical crimson to prototypical red is one indicator, but does not suffice by itself, since it does not discriminate between “All crimson things are red” (true) and “All red things are crimson” (false). Various cognitive mechanisms can be postulated to do the job. We need not fill in the details, since for present purposes what matters is the overall picture.¹⁴

This is, to put it kindly, somewhat cavalier. What is “the overall picture”? It is no good just gesturing airily at unstated “cognitive mechanisms”. We will not regard Norman as knowledgeable about (*Crimson*) unless he refuses the converse generalisation. Since Williamson's description of Norman's cognitive processing underdetermines the latter's ability to make this discrimination, it has to be at best incomplete. But then no relevant analogy with the *Who's Who* case has yet been made out – for until the missing details, are filled in, we don't know whether or not they might be such as to cause the analogy to break down.

In general, it seems that before a process of visualisation of a crimson prototype and generalisation from its character can lead to knowledge about relations of colour inclusion, the agent has to have some anterior conception of the spread of the band of shades of which the visualised shade is typical, and the spread of the putative determinable colour concerned. The trouble is that if Norman has that anterior conception, it's not clear what work the episode in the visual imagination plays in securing the knowledge. If he knows already how far the crimson band spreads, as it were, and, for its part, how far the red band spreads, then he already knows all he needs to know to judge that all crimson things are red while, conversely, not all red things are crimson. The interesting question will then concern the source of those two pieces of knowledge. But whatever it may be, it looks as though that is where Norman's knowledge of colour inclusions and exclusions is grounded. What if any part may be played by the kind of episode in visual imagination—rather than, say, explicit training in the relevant linguistic boundaries—that Williamson describes, is, as far as his discussion succeeds in taking matters, quite moot.

In summary: Williamson's idea was that skills of imagination, honed by perceptual experience, can play the same kind of role in the acquisition of knowledge of propositions as diverse as “All crimson things are red” and “All recent volumes of *Who's Who* are red”. But the effect of the foregoing is that his discussions have left it quite unclear what role the visual imagination *does* play in the epistemology of the former. Accordingly, no analogy suitable for his purpose has yet been made out.

¹⁴ Williamson (2013: 296).

§7 *A better example for Williamson's purpose?*

Arguably, though, this is a drawback of the particular example that Williamson selected and can be finessed. Let's switch to one that we offer as more plausible for his purposes. Consider the following:

(*Square*) All squares are diamonds.

In this case there is a much more salient role that the visual imagination can play in inducing conviction, which is relatively easy to describe. Suppose that Norman has acquired the concepts *square* and *diamond* in the kind of ostensive way that Williamson suggests and now considers for the first time whether (*Square*) is true. He visually imagines a stereotypical square and then *realigns it in the imagination* by rotating it through 45 degrees in such a way that it then impresses as diamond-shaped. Satisfied with this, he assents to (*Square*).

Questions still remain to be answered about this process, of course. There is still the problem of the erroneous converse generalisation — Norman might as well have started with a visually imagined picture of a square resting, so to say, on one of its corners and then, realigning it in the imagination to present as a square, have concluded that all diamonds are squares. So something needs to be said about the original process to explain why it is productive of knowledge although the latter would not be. There is still an issue of incompleteness, therefore, in the description. And there is also the problem of generalisation – why does Norman get to be justifiably certain that *all* squares are diamonds, rather than merely that this particular square, re-configured in his imagination, is a diamond? But at least it is clearer in this case how a process taking place in the imagination intelligibly prompts the belief that is formed.

§8 *Questioning the Core Analogy (II)— Lemmas and Props*

So, is Norman's process with (*Square*) epistemically relevantly similar to the process that Williamson describes when Norman is asked about the colour of recent volumes of *Who's Who*?

There is an evident disanalogy. We can bring it out by focusing on the structure of the epistemic justification supporting Norman's judgment. Compare the following triads:

<i>Who's Who</i>	i	Norman's visual imagination of a volume of <i>Who's Who</i>
	ii	Norman's judgment: all recent volumes of <i>Who's Who</i> are red
	iii	The publishers of <i>Who's Who</i> use the same colour for all the volumes.
<i>Square</i>	i	Norman's realignment of a visually imagined square through 45 degrees
	ii	Norman's judgment: all squares are diamonds
	iii	The visualised square is stereotypical and any square may “in principle” be reorientated in that way.

In each case proposition iii is an *authenticity condition* on the relevant warrant: a cogniser who doubted it could not rationally believe ii just on the basis of i.¹⁵

While every epistemic warrant relies on a number of authenticity conditions, what makes the two we conditions we have singled out especially interesting is that they are crucial in underwriting the generality involved in Norman's judgment ii. They do so in the following way. If Norman doubted *Who's Who* iii, he could not rationally move to, or justify, the proposition (*Who's Who*) on the basis of the imaginative episode i that Williamson described. And if Norman doubted *Square* iii, he could not rationally move to, or justify, the proposition (*Square*) on the basis of the imaginative episode *Square* i.

Given that we are discussing the warrants for two general claims, and that the two authenticity conditions individuated underwrite the reliability of the respective generalization steps, one may expect that propositions iii are properly regarded as *lemmas* in the epistemic architecture of Norman's two judgements. That is, one may expect that they are propositions whose truth Norman needs some assurance about before he can justifiably take the imaginative processing described in the two i-components as grounding the generalisations those judgements make. Lemmas contrast with *props*, authenticity conditions for which independent reassurance need *not* to be demanded before one can warrantably make the relevant judgment. For example, one might expect that a condition stating that Norman's cognitive abilities are not impaired at the moment in which he makes the judgments would be a mere prop for Norman's warrants. Suppose then, if only for the sake of the argument, that the two iii propositions are indeed lemmas.

Now consider an evident difference between the two authenticity conditions: removing a doubt about *Who's Who* iii would require routine empirical investigation — a call to the publishers, a trip to the public library, other people's testimony, or whatever. By contrast, the assurance – if needed – that the visualised square is stereotypical and the process or reorientation always possible, at least “in principle”, requires no such investigation. On the contrary, by bottom-up standards, and although its grounds may seem mysterious in the present context, it's an assurance that can be accomplished in stride and, intuitively, a priori— we find the point *obvious on reflection*.

Thus, if the status of the two iii- propositions is that of lemmas, the Williamsonian analogy breaks down. Part of the process of Norman's rationally persuading himself, by the relevant phenomenological exercises, that the two generalisations hold will be the attainment of independent assurances that the two lemmas hold, and at that point the claimed analogy will break down. On this account the appearance of analogy is indeed sustained only by the incompleteness of Williamson's style of description of the cognitive processing involved in the two cases. Supplementary information is needed before Norman can achieve knowledge of (*Square*) and of (*Who's Who*) in the kind of way described; and this supplementary information is empirical in the one case, but a priori in the other.

It is true, of course, that this way of making a distinction *relies on* the notion of apriority — on the claim that the needed lemma is a priori in the one case but in need of empirical confirmation in the other. Against a convinced sceptic about the intelligibility of any such distinction it is therefore dialectically ineffective. But Williamson, remember, is not

¹⁵ For the terminology of “authenticity conditions” and some discussion of the notion, see Wright (2007: 29-30, 2014: 214ff), and Wright (2004: 191), where the notion was prefigured under the name of “presuppositions of cognitive projects”. Authenticity conditions coincide at least in extension with what Pryor (2012: 298) calls ‘anti-underminers’.

doubting that we do have a “bottom-up” distinction. So he can have no objection to a critic's appealing to it in order to explain a difference between cases that he is arguing are essentially similar.

But *are* the iii-propositions correctly regarded as lemmas? Not all conditions a doubt about which would undermine justification of a certain kind for a claim are conditions for which independent reassurance needs to be demanded before one has such justification. Any coherent epistemology has to reckon with conditions of this type — has to be, in contemporary terminology, *liberal* with respect to some such conditions.¹⁶ The hard question is which are the conditions – the potential underminers – which serve merely as background *props* and about which one can take a liberal view, and which are those which function as lemmas and for which the opposing *conservative* view is appropriate: the view that, in these cases at least, independent assurance must be part of any rational conviction about the truth of the target proposition acquired on the basis of a certain type of ground? In these terms, the analogy-breaker just outlined rests on a conservative view of the two iii- conditions. What if Williamson were to counter by taking a liberal view—by proposing that the conditions concerned are merely props?

Well, the question, when in general to be liberal about authenticity conditions, or anti-underminers, and when to be conservative, is one we think that no one at present knows how to answer. But there is a case for thinking that a liberal view would be quite inappropriate in the case of (*Who's Who*). For suppose that, instead of a play within the visual imagination, Norman is presented with an actual token of a recent volume of *Who's Who*, and an actual drawing of a square on a piece of paper. Surely this effects no relevant change in his epistemic position with respect to (*Square*) and (*Who's Who*). He holds the paper in his hand, rotates it appropriately, and concludes as before that all squares are diamonds. He looks at the volume of *Who's Who* and concludes as before that all recent volumes of *Who's Who* are red. In the latter case we surely *must* require that he have some independent reason to believe that the volume he is looking at is typical before he can be regarded as justifiably generalising. In the former case, it seems there is vanishingly small room for rational doubt that the ability of the square to present the two different aspects, depending on its orientation to the observer, is essential to it and may be safely attributed to squares in general. However that may be, epistemic responsibility seems quite clearly to require conservatism about the iii-condition in the *perceived* volume of *Who's Who* case. And that is enough to set up a dilemma. If we take the view that Norman does strictly need independent assurance that the presentation of distinct *gestalten* by the square that he has physically reoriented can be expected to be sustained by squares in general, still it doesn't seem as though that reassurance needs to be secured empirically. But if we take the view that he needs no such independent assurance – that he may legitimately simply generalise from his experience with the paper square that he manipulates – then that is also a crucial distinction between the two cases.

In the concrete case, then, the analogy breaks down, one way or the other. Why would—how *could*—it make a relevant difference, then, if the relevant tokens are presented not in external experience, but in imagination? Given Williamson's reliance on the analogy

¹⁶ Cf. Pryor (2004: 356).

between online (perception) and offline (imagination) ways of coming to know, this would seem to be an especially pointed question.¹⁷

§9 Coda—*Learning from Imagination.*

As we have emphasised, we regard it as important for the interest of Williamson’s sceptical project that it avoid dependence upon externalist preconceptions about knowledge. In setting up the analogy, Williamson takes it for granted that the kind of imaginative processes he describes — processes involving experience in the more-than-merely-enabling but non-evidential way — are indeed generally reliable across a range of bottom-up a priori and bottom-up a posteriori cases of which he intends his examples to be typical. Let that be so. It is notable that reliabilism would provide a framework in which, for the purposes of determining whether a certain procedure was generally knowledge-productive, the distinction between lemmas and props would be degraded: both alike would merely be conditions on the reliability of the procedure. It is tempting to speculate that the considerations gestured at in the preceding section are overlooked in Williamson’s account precisely because he is, implicitly, viewing Norman’s envisaged processes through reliabilist (or other externalist) lenses.

Be that as it may, the question remains: what is the source of the reliability in the two kinds of case? Is it relevantly similar?

Suppose Norman runs through the process that Williamson describes, and comes to the opinion, potentially knowledgeable in Williamson’s view, that all recent copies of *Who’s Who* are red. Suppose we are somewhat incredulous, and ask him whether he remembers seeing any copies recently, and suppose, in accordance with Williamson’s specification of the case, that he replies that he recollects no particular such episode. He insists that all he did was to visualise a specimen. What reason is there why we, or he, should place any credence in his flight of imagination?

Williamson’s answer may be found in §2 of his later paper, “Knowing by Imagining”.¹⁸ There he speculates that the imagination originally evolved as a useful means to alert our remote ancestors to potential dangers and opportunities, and argues that in order to serve that purpose well, it had to be both selective and reality-oriented. Supposing that such is the way in which it evolved, it is to be expected that imagination, in a wide range of circumstances, will be reliable enough to confer knowledge.

“But what”, we may ask, “about the evident possibility of wild and fanciful imaginative exercises of a broadly similar phenomenology to the epistemically worthwhile ones? What is to stop Norman freely imagining a pink polka-dotted edition of *Who’s Who*, unlike any ever produced, and then generalising from that?” Again, the later paper might offer an answer. Williamson proposes a distinction between *voluntary* and *involuntary* modes in which the imagination operates (§2), and appears to suggest, very roughly, that the knowledge-conferring episodes of imagination—as opposed to the fanciful ones—are those

¹⁷ That the analogy with ordinary cases of perception drives Williamson’s account of the epistemic role of the imagination is even clearer in his (2016), where he explicitly claims (p. 118) that the online and offline processes take the same input and deliver the same output by the same means, and differ only with respect to the sources of input (perception in one case, imagination in the other). Williamson stretches the analogy to the point of saying that scepticism about the epistemic role of imagination leads to scepticism about the epistemic role of online faculties (119).

¹⁸ Williamson (2016)

in which, while the scenario from which the imaginative exercise begins (in our examples, the visualization of a crimson patch, or square, or of a recent copy of *Who's Who*) is enacted voluntarily, the conclusion of the imaginative exercise (the recognition of the redness of the crimson patch and of the volume, or the diamond-shape of the square) unfolds *involuntarily* (§3). According to Williamson's admittedly speculative story, "left to itself, the imagination develops the scenario in a reality oriented way, by default" (2016: 116).

A critic may be inclined to put pressure on this play with voluntary and involuntary aspects of imaginative exercises. How, for example, does it supply the materials to discount the epistemic significance of the kind of case when, in a daydream say, a vivid imaginary scene pops up, as it seems, wholly involuntarily? But we shall not press such issues in this already lengthy review. Williamson's evolutionary hypothesis is, after all, not implausible and there are, surely, many kinds of example where reliance on imagination does characteristically generate valuable information. The examples Williamson gives in his (2016) include knowledge of what would happen in hypothetical circumstances (§3), and some kinds of knowledge of the future, and of the past (§4). It is notable, however, that this later paper does not explicitly discuss the role of imagination in bottom-up a priori cases or attempt to apply the evolutionary hypothesis to account for its efficacy there. If it had, that would surely have marked the demise of any ambition that Williamson's core thesis—that the contrast between a priori and a posteriori marks no deep or important distinction among different ways of knowing—should not rely upon externalist assumptions about knowledge. For the evolutionary story is externalist through-and-through.

There is, though, a salient and significant contrast in the grounds of the reliability of the imaginative process in the two kinds of case we have been concerned with. If we ask Norman why he believes that all squares are diamonds, and he replies that he imagined one and reoriented it in the imagination, we don't in general feel any need whatever for an independent check on the credentials of this kind of process. Indeed, to feel the want of one would seem to betray a misunderstanding of what Norman has actually done. When the imaginative process is fully described, we too will find it fully convincing, for Norman's routine with the imaginary square does not, as it were, bear a merely *contingent and circumstantial* relation to the truth of the proposition that it suggests. Indeed not only do we not, in basic a priori cases, require an independent check on the reliability of the processing in this way; it's not clear in general how, if we *did* require it, the check might be accomplished. Perhaps the proposition in question might be one that allowed of independent proof. But that will not be the general run of cases. Alternatively we might check out the generalisation inductively; but the utterly convincing process in the imagination that underwrites the generalisation that all squares are diamonds creates a barrier to the intelligibility of the idea that an inductive check on the generalisation might disclose counterexamples.

None of that is true of the imaginative episode that convinces Norman of (*Who's Who*). Even if we grant that Norman's phenomenological routine is essentially the same for both (*Square*) and (*Who's Who*), the reliability of the connection between what he does and the belief that he forms seems to have a quite different status in the two kinds of case: a difference which we can loosely summarize by saying that, in the (*Who's Who*) case, any correlation between the process and the truth of the product is indeed contingent and circumstantial—exactly as it is represented as being by Williamson's evolutionary speculations—whereas in the case of (*Square*) it is not. That of course is exactly the

difference we would expect if we were to regard what Norman does in the case of (*Square*) as a kind of *informal proof* of that proposition, and what he does in the case of (*Who's Who*) as a merely experimental operation. To be clear: the reason why we should not think of the belief-forming process in (*Who's Who*) as an informal proof is that the authenticity condition that concerns the generality of the putatively warranted claim is such that (a) it can only be checked empirically, and (ii) it seems to play the role of a lemma rather than that of a prop (so that the empirical check is needed before the relevant belief may count as warranted).

It's natural to conjecture, therefore, that the real difference of importance marked by the a priori/a posteriori distinction as drawn bottom-up is to be located at this point. There may well be methods of corroboration that one can execute in the mind's eye, as it were, whose products are, by bottom-up standards, ordinary a posteriori propositions. But their connection with the truth of those products will be contingent, circumstantial, and open to experimental corroboration. Not so with the relationship between a proof, even a "baby" proof, constituted by a simple informal routine in the imagination, and the truth of what is proved. In such cases, the relevant belief-forming processes are such that, *if executed correctly*, they guarantee a true outcome.

In his notes posthumously published as the *Remarks on the Foundations of Mathematics*, Wittgenstein repeatedly returns to the question, what is the difference between a proof and an experiment? And his answer, roughly, is that the difference resides in the *attitude* we take to the relevant process. A majority, perhaps, will not much care for that answer, but the importance of the question for epistemology generally can hardly be doubted. It is an insight of Williamson's discussion that some processes executed "in the head" can, in the right circumstances, provide experimental corroboration of ordinary empirical propositions concerning external matters. But that reflection does nothing to call into question the substance of the contrast between proving a proposition and experimentally confirming it. Perhaps a reminder is timely that it is that contrast that the traditional terminology of *a priori* and *a posteriori* methods of acquiring knowledge gestures at.¹⁹

Works cited

Casullo, A. Ms. 'A defence of the Significance of the A Priori – A Posteriori Distinction.' This Volume.

Hawthorne, J. 2007. 'A Priority and Externalism.' In S.C. Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology*. Oxford: Oxford University Press, 201-218.

Jenkins C. and Kasaki M. 2015. 'The Traditional Conception of the A Priori'. *Synthese*, 192: 2725-2746.

Pryor, J. 2004. 'What's Wrong with Moore's Argument?'. *Philosophical Issues*, 14, 349-378.

¹⁹ Work on this article has received funding from the European Union's Horizon 2020 Research and Innovation programme under Grant Agreement no. 675415.

Pryor, J. 2012. 'When Warrant Transmits.' In A. Coliva (ed.), *Mind, Meaning, and Knowledge. Themes from the Philosophy of Crispin Wright*. Oxford: Oxford University Press, 269-303.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, T. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.

Williamson, T. 2013. 'How Deep is the Distinction between A Priori and A Posteriori Knowledge?' In Casullo, A. and Thurow, J. (eds.), *The A Priori in Philosophy*. Oxford: Oxford University Press, 291-312.

Williamson, T. 2016. 'Knowing by Imagining'. In A. Kind and P. Kung (eds.), *Knowledge Through Imagination*. Oxford: Oxford University Press, 113-123.

Wright, C. 2004. 'Warrant for Nothing (and Foundations for Free?).' *Proceedings of the Aristotelian Society*. Supp. Vol. LXXVIII, 167-212.

Wright, C. 2007. 'The Perils of Dogmatism', in S. Nuccetelli and H. Seay (eds.), *Themes from G. E. Moore: New Essays in Epistemology and Ethics*. Oxford; Oxford University Press, 25-48.

Wright, C. 2014, 'On Epistemic Entitlement (II): Welfare State Epistemology', in D. Dodd and E. Zardini (eds.), *Scepticism and Perceptual Justification*. Oxford; Oxford University press, 213-247.