Is the *A Priori/A Posteriori* Distinction Superficial?

[penultimate version of contribution to Blake Roeber, Ernest Sosa, Matthias Steup, and John Turri (eds.), *Contemporary Debates in Epistemology*, 3rd edition, Wiley-Blackwell, 2024]

Timothy Williamson

Introduction

Using my eyes, I can see that a computer screen is there. With my eyes shut, I can see that the conclusion of a syllogism follows from the premises. The first ‘see’ has a specifically *visual* sense, the second an *intellectual* sense, but the dead visual metaphor indicates some similarity between the two cases. At the very least, in one I come to *know* that a computer screen is there, in the other I come to *know* that the conclusion follows from the premises, in the same unspecific sense of ‘know’. Yet many philosophers claim a deep difference between the two cases: whereas I know that a computer screen is there *a posteriori*, I know that the conclusion follows from the premises *a priori*.[1] On the standard philosophical use of ‘*a priori*’ and ‘*a posteriori*’, those classifications may be correct. But that does not show that the difference goes deep. A toy analogy: we correctly classify some bicycles as *red*, and others as *not red*, but that does not show that there is a deep difference between red bicycles and non-red bicycles. More seriously: an attempt to define ‘races’ by artificial DNA-based criteria might somehow succeed in partitioning all humans into a few mutually exclusive, jointly exhaustive classes, yet reveal nothing of deep theoretical significance.

The traditional one-liner about the difference between the *a priori* and the *a posteriori*: *a priori* knowledge is *independent of experience*; *a posteriori* knowledge *depends on experience*. Of course, in the usual sense of ‘experience’, both seeing that a computer screen is there and seeing that the conclusion follows from the premises are, in their own small way, experiences. Without them, I would not have known those truths, or at least would not have come to know them in the way I actually did. Thus friends of the *a priori/a posteriori* distinction (henceforth, ‘the Distinction’) must do theoretical work to clarify what they mean by ‘experience’ and ‘depend’, or else explain the Distinction in other terms, to make it deliver the outcomes they want. I will not discuss such attempts in detail here. Instead, I will focus on a standard paradigm of *a priori* knowledge, and show how similar it is epistemologically to some paradigms of *a posteriori* knowledge. Although one can still draw a theoretical line between the two classes, it fails to mark a boundary of much epistemological significance; it does not cut at a cognitive joint.

The point is not that the Distinction has borderline cases. Most useful distinctions have those. What shows the distinction between red bicycles and non-red bicycles to be superficial is not a bicycle that is neither clearly red nor clearly non-red, but rather making a clearly red bicycle clearly non-red just by painting it green, changing nothing important.

Analogously, I will focus on similarities between clear cases of *a priori* knowledge and clear cases of *a posteriori* knowledge.

Example

A central kind of *a priori* knowledge is knowledge based on logical or mathematical proof. Unless such cases normally fall under the category of *a priori* knowledge, it risks being too sparsely inhabited, too marginal, to matter much. The Distinction's true friends should be willing to maintain that proof-based knowledge is normally *a priori*.

Here is an example. Late one night, a mathematician has an idea for a proof of a new theorem, the mathematical statement M. She hurriedly writes it down. The argument is long and complex. She is unsure whether it is correct; she does not yet know the truth of the theorem. The putative proof needs to be checked. Having written it down, exhausted, she decides to leave the work of checking it until she is fresh after a night's sleep. The next morning, she painstakingly goes through the proof Π, checking each step, and verifies its correctness by normal mathematical standards. She now knows that Π is a (correct) proof of M, and thereby knows M itself. This is an ordinary instance of knowledge based on mathematical proof. As just explained, friends of the Distinction will count her as knowing M *a priori*. Moreover, she knows M *by* knowing that Π is a proof of M. In some simple cases, one might know a mathematical truth just by proving it, without ever thinking *about* the proof, but in trickier cases—like this one—reflection on the proof itself is epistemologically crucial. For the mathematician to know M *a priori* by knowing *a posteriori* that Π is a proof of M would be bizarre, and already evidence that the Distinction cuts at no joint. Anyway, that Π is a proof of M is itself a mathematical fact; proof theory is a branch of mathematical logic. I will assume that friends of the Distinction count the mathematician as knowing *a priori* that Π is a proof of M. I will focus on her *a priori* knowledge of the proof.

What does the mathematician do when she checks the proof? For clarity, we may assume that her checking is unusually minute and meticulous, by ordinary mathematical standards. For instance, she looks at a passage written in her notebook, and verifies that it constitutes a genuine instance of modus ponens, where she makes a transition from premises 'If A, C' and 'A' to a conclusion 'C' (which she knows to be truth-preserving). She can do this because she can recognize an instance of modus ponens when she sees one, in her own language (she would not recognize an instance in an unfamiliar language, with a word she did not understand instead of 'if'). Her ability to recognize such an instance of modus ponens when she sees one is a *perceptual recognitional capacity*. It involves a form of pattern recognition, like a chess grandmaster's recognition of an abstract pattern in the layout of the pieces on a board. Probably, she can also recognize an instance of modus ponens when she hears one, though that is not quite the same skill (it depends on aural memory, especially when intervening material separates the premises and conclusion); the chess grandmaster may be unable to recognize the pattern of pieces by touch (a black pawn and a white pawn feel the same). Naturally, a visual-recognitional capacity for modus ponens is *selective*; it ignores many aspects of what is seen. But the same holds of visual-recognitional capacities for shapes and colours too.

The mathematician's visual recognitional capacity enables her to see that *an instance of modus ponens is written there*. Since her attention is less on where it is written than on whether it is a genuine instance of modus ponens, we might more relevantly describe her as seeing that *this is an instance of modus ponens*. But her knowledge that *an instance of modus ponens is written there* and her knowledge that *this is an instance of modus ponens* derive from the same underlying visual recognitional capacity. In the latter knowledge, the demonstrative '*this*' is perceptual: its reference is fixed by her visual attention, just as the reference of the perceptual demonstrative '*there*' in the former knowledge is fixed by her visual attention to a place on the page. Such recognitional capacities are likely to have been learnt and calibrated in sense perception, perhaps by trial and error—through experience, one might say—and so owe much of their reliability to sense perception, however they are subsequently applied. For logically significant patterns, such as modus ponens, the learning may have been fused with language learning, of logically significant words such as 'if'.

Uncontroversially, the mathematician cannot know *a priori* that *an instance of modus ponens is written there*; she must know it *a posteriori*. Nevertheless, might her knowledge that *this is an instance of modus ponens* somehow count as *a priori*, even though the reference of the demonstrative is fixed by her visual attention? Perhaps, but that would support my claim that paradigms of *a posteriori* knowledge are very similar to clear cases of *a priori* knowledge.

Friends of the Distinction may complain that my emphasis on the written form of the proof is a red herring, because in principle the mathematician could verify the whole proof in her head. Is that a good response?

The qualifier 'in principle' is a euphemism for a huge idealization. Many proofs in mathematics are far too long and complex for anyone to hold in their head as a whole. Even if each individual step can be entertained in the head, verifying a proof also involves checking that the individual steps all fit together properly into a proof of the required conclusion. It is unclear how one could do that without some lasting record of the proof's individual components, such as writing makes available. For long proofs, human memory is insufficiently reliable.

The epistemology of vision is sometimes claimed to be irrelevant to seeing a proof, since hallucinating it would do just as well. That might have some plausibility for individual steps, but not for the proof as whole. If you had just gone through a written proof, making painstaking checks, and then were told that you had been hallucinating half the time, you would be in no position to say 'It doesn't matter, I still have a valid proof'. You would not know whether all the individual steps fitted together properly into a proof of the conclusion.

Of course, we can schematically *imagine* super-humans whose consciousness is vastly more capacious than our own, enabling them to take in and verify the proof in a single act of consciousness. But if our own knowledge of the theorem depends on sense perception, so does any knowledge we may have that a superhuman could know the theorem independently of sense perception.

Another consideration is that mathematics is a collective enterprise. Proofs must be checkable by the community of mathematicians, just as experiments in natural science must be repeatable. One obvious mechanism for that is the refereeing process for articles to be published in mathematical journals. A mathematician's solemn promise that he has a clear proof in his head is not enough. A publicly available proof, checked by other mathematicians or a computerised proof-assistant, is the epistemological gold standard.

In those respects—surveyability and public checkability—a written proof beats one in the head. In other respects, a proof in the head is quite similar to one on paper. Visualizing an instance of modus ponens is the offline version of the online process of seeing the instance; they differ primarily in the source of the inputs. One uses a visual recognitional capacity online to recognize instances of modus ponens in what one sees; one can reuse the same recognitional capacity offline to recognize instances of modus ponens in what one imagines. Similarly, expert chess players may use the same capacity for pattern recognition when they imagine a configuration on the board as when they see it on the board. Of course, in mentally rehearsing a proof, one is not confined to imagining *formulas* (as for modus ponens); one may also imagine *diagrams*. But an imagined diagram plays much the same role as one seen on a board, though many of us find the one drawn on the board clearer and more stable. The same underlying visual recognitional capacities are in play, online or offline (for related discussion see Carruthers 2015). The main advantages of doing mathematics in one's head, when one can, are just that it is comparatively quick and convenient.

In short, to switch from going through a proof on paper to going through it in one's head does not avoid reliance on sensory-perceptual skills. It is just a switch from relying on them online to relying on them offline. If that switch made the difference between *a posteriori* and *a priori* knowledge, the difference would be rather superficial, and mainly to the advantage of the *a posteriori* side.

Friends of the Distinction may try a different line, objecting that such arguments go wrong because they focus on the perceived *forms* of representations instead of the *contents* which they represent: a proof's written form merely enables access to the *real* proof, a more abstract intellectual structure.

The trouble with that objection is that it neglects what formal proofs are *for*. A generally accepted requirement for a proof system to count as *formal* is that there should in principle be a mechanical procedure for determining of any given proof-candidate (a sequence or array of representations) whether it constitutes a correct proof in the system. Any such procedure can be implemented on a computer, given appropriate ways of scanning and coding proof-candidates (like Gödel numbering). Although it may be controversial whether the system's basic axioms and rules of inference are sound on their intended interpretation, there should in principle be no *additional* controversy about the correctness of individual proofs in the system. Such a mechanical decision procedure for proofhood has to operate on formal aspects of proofs, not on their content.

For example, suppose that the computer has to determine whether X, Y, and Z constitute an instance of modus ponens, with X the major premise, Y the minor premise, and Z the conclusion. In effect, it must determine whether X is the conditional with antecedent Y and consequent Z. For a language whose formulas are sequences of symbols from a finite alphabet, with conditional formulas written left-to-right (left parenthesis, antecedent, conditional symbol $\rightarrow$, consequent, right parenthesis), the question boils down to whether X is the same formula as $(Y \rightarrow Z)$. That question is easy to answer for a computer attached to a suitable scanning device, with recognitional capacities for each symbol of the alphabet. By contrast, if the premises and conclusions of arguments were pure abstract contents, with no such quasi-syntactic structure, it is quite unclear what could be meant by a 'mechanical test' for instances of modus ponens. Thus treating perceptible form as epistemologically irrelevant risks throwing away the very feature on which the epistemic value of formal proof depends.

More generally, when one looks at a standard presentation of the rules in a system of natural deduction for a logic, one sees the introduction and elimination rules for the various connectives displayed as abstract visual patterns (modus ponens is the elimination rule for the conditional). Learning the rules partly consists in acquiring visual recognitional capacities for those patterns, which one can use both online in perception and offline in imagination. Such natural deduction rules are sound and complete for first-order logic and approximate well to the background logic of most ordinary mathematical reasoning.

Of course, most proofs in working mathematics are not purely formal. Constructing and understanding them still requires pattern-recognition, but the patterns are typically 'macroscopic' and specific to the relevant subfield (appeals to Church's Thesis in recursive function theory are an extreme case in point) rather than 'microscopic' and general (such as modus ponens). Numerous short-cuts save time and space and prevent clogging detail from obscuring key ideas. Such proofs are written to be understood and checked by experts in the subfield, and are too informal to be checked by a computer; they may never get fully formalized. But we have no reason to expect such macroscopic pattern recognition to depend any less on forms of representation than does microscopic pattern recognition; it just depends on forms of representation perspicuous for macroscopic patterns. Moreover, when doubts about a proof intensify, it may require scrutiny at increasingly microscopic levels.

Mathematics faces a growing problem with proofs so long and intricate that even leading experts in the subfield cannot be confident whether the proofs are correct just on the basis of checking them 'by hand', let alone 'by brain'. In response, there is a growing trend towards computer-assisted proofs, where human mathematicians interact with computer assistant programs: the human enters raw definitions and sub-proofs, the program gives prompts whenever it finds a definition unclear or a step unobvious, the human provides clarifications or intermediate steps, and so on. After several months of such work, a very difficult proof-sketch by the Fields medallist Peter Scholze of an important new theorem was recently verified (Castelvecchi 2021). Thus formal standards of proof continue to play a key epistemological role in contemporary mathematics. Naturally, mathematicians hope to simplify and streamline complex proofs to make them more humanly intelligible, but it is an open question how far such hopes can be realized. In any case, if one treats the overt form of mathematical representations as epistemologically irrelevant, it is hard to explain the role of formal proof as a standard to which ordinary mathematical proofs are in principle held. How is a computer supposed to check a proof consisting of pure contents with no quasi-syntactic form? In short, form is crucial to mathematical proofs, from the simplest and most elementary to the most complex and advanced.

The illusion that overt syntactic form is irrelevant to an underlying 'pure' proof may result from the arbitrariness of basic logical and mathematical symbols. For instance, what symbol one uses for the conditional matters little. Although '$\rightarrow$' seems rather perspicuous for its meaning, most ordinary mathematical proofs use a natural language conditional; for mathematical purposes, English 'if' is neither better nor worse than Italian 'se'. Let us extend the term 'language' to whatever system of representation (formal or natural language, mathematical notation, diagrams, …) is used in a proof. Since mathematical proofs can so often be translated from one language into another without mathematical loss or gain, one might get the impression that the language is irrelevant to an underlying language-free proof. But that would be a fallacy. For although no *particular* language is needed for a given type of proof, it does not follow that such a proof can work *with no language at all*. Analogously, if a story can be told in any natural language, it does not follow that it can be told without use of language. Thus, when we consider the epistemology of proof, we must not treat the language as something to be *factored out*. Its perceptible or imaginable presence is crucial, even though many different languages will do equally well. Without some language or other, we cannot even ask mathematical questions, let alone reason our way to an answer.[2]

One manifestation of the centrality of linguistic form to proof is in the role of *free variables*, ubiquitous in mathematics. For example: where 'x' and 'y' are distinct variables, we cannot derive C from Fx and $Fy \rightarrow C$ by modus ponens, even though the variables do not differ in non-linguistic content. In assessing this inference, exhortations to attend to its content, not its form, would merely confuse the issue.

The centrality of representational form is not confined to proof *checking*. It is central to proof *construction* too. The idea for a proof often comes itself from recognizing a pattern, in formulas, a diagram, or whatever, and the proof may be developed by manipulating such patterns, playing with them. Anyway, proof construction is not a separate process from proof checking. As the overall proof is constructed, the mathematician will be continually checking proofs of lemmas, steps towards the final destination.

I have *not* argued that logical and mathematical proofs yield only *a posteriori* knowledge. If some philosophers want to stipulate that knowledge based on logical or mathematical proof is simply paradigmatic of *a priori* knowledge, let them. For all that, the knowledge was obtained through exercise of the same recognitional capacities through which we can also obtain paradigmatically *a posteriori* knowledge.

Evolutionary considerations

We should not be surprised that paradigms of *a priori* knowledge involve ordinary perceptual-recognitional capacities, used in special ways, online or offline. What else would one expect on evolutionary grounds? How much does knowledge of formal logic, mathematics, or philosophy increase one's chances of surviving to breed successfully and pass on one's genes? A special cognitive capacity dedicated to the *a priori*—for instance, a *sui generis* capacity to have intellectual seemings of necessary truth—would not pay its evolutionary way. Extra brainpower costs energy. A far more plausible explanation of the human capacity for *a priori* knowledge of formal logic, mathematics, and philosophy is that we do it by ingeniously applying our ordinary cognitive capacities for purposes they never evolved to serve, for instance in ways just sketched.

Of course, such evolutionary considerations are not decisive. Not everything has an evolutionary function. There are *spandrels*, mere by-products of something else which did serve an evolutionary function. Perhaps a special cognitive faculty dedicated to the *a priori* came along for the ride with the easiest evolutionary means to develop some fitnessincreasing feature. But the spandrel hypothesis is at a disadvantage, compared to an account like that above. We do better to explain a feature of human cognition by its evolutionary function than by postulating it to be a by-product of something else in some unspecified way.

Clearly, many human cognitive achievements are to be explained culturally rather than biologically. Such explanations can still avoid evolutionary implausibility, for example by showing how human communities found and developed new ways to use their basic cognitive capacities.

Knowledge by proof is a case in point. The traditional medium of proof in logic and mathematics is distinctively human: language, aided by diagrams and symbols. Properly to keep track of complex proofs, writing is needed. Formal proof may go back no further than Euclid, though less formal proofs came earlier. Geometry itself originated with techniques for solving practical problems: if you can measure the sides of a field, and thence calculate its area, you can estimate the crop yield. Until the nineteenth century, the intended subject matter of geometry was the structure of physical space; that was why the consistency of non-Euclidean geometries felt so threatening. Of course, as a branch of contemporary mathematics, geometry studies abstract structures, whose relation to physical space it leaves to physics. But mathematicians still use what are in effect spatial metaphors as powerful means for understanding abstract structures. One sign of that is the role of diagrams in mathematical proof, highly valued by most mathematicians, which recruits their general abilities in spatial reasoning and manipulation for more abstract purposes (see for example De Toffoli 2017). In doing so, they have learnt how to avoid reading more into spatial diagrams than they are meant to represent. The development of such mathematics by intelligent language-using creatures with evolved capacities for spatial reasoning and manipulation is not utterly mysterious.

One might say: formal proofs in logic stand to informal verbal reasoning as formal proofs in geometry stand to informal spatial reasoning, with similarities in their historical development. However, even logic has a spatial aspect, introduced both by the need for writing to keep control of complex proofs and by the use of diagrams in more mathematically sophisticated logic.

Our ability to reason in our heads as well as on sand, paper, or board is part of a general human capacity to work offline, in the imagination, as well as online, interacting with what we sensorily perceive. Our cognitive capacities can typically be used both online and offline. Our capacity to work offline has significant practical value. In decision-making, we have to judge what will or may happen *if* we select a particular option: the natural way for us to determine that is by *imagining* selecting that option and realistically thinking through the consequences. Usually, such consequences are causal rather than logical, and by no means *a priori*. But by suitably filtering the background knowledge and belief we use in developing the initial supposition, we can in effect restrict ourselves to logical or mathematical consequences, as in a proof. This is a limiting case of our general capacity for reasoning under a supposition.

These remarks are obviously just a beginning. But they suggest that there is no need to invoke a cognitive *deus ex machina* to provide special ingredients to cook up paradigms of *a priori* knowledge; the usual ingredients suffice.

These challenges may tempt some simply to *identify a priori* knowledge with innate knowledge and *a posteriori* knowledge with acquired knowledge, assuming the scientific respectability of the innate/acquired distinction. However, that distinction is itself contested, on whether it can be defined scientifically and, if so, how (Griffiths 2020). Even if we put those concerns aside, it is not clear that the innate/acquired distinction will do what the Distinction is intended to do. For example, suppose that humans turn out to have an innate fear of snakes, for good evolutionary reasons (Kawai 2019). We may have innate knowledge that some snakes are dangerous. But do we really know *a priori* that some snakes are dangerous? More generally, on such views, the status of innate knowledge as *knowledge* depends on its evolutionary origins. In effect, it is a form of learning from interaction with the environment, at the level of the species rather than the individual. Even if that is not exactly the same as individual learning from perceptual interaction with the environment, friends of the Distinction may well feel that it is closer to paradigms of *a posteriori* knowledge than to *a priori* knowledge. Conversely, most mathematical knowledge is obviously not innate. This is clear for non-elementary mathematics, but even at the most elementary level, no recognitional capacity for modus ponens for 'if' is innate, since knowledge of English is not innate. According to Jerry Fodor (1975), we have a language of thought (Mentalese); there might even be an innate computational capacity to identify and accept instances of modus ponens for a conditional device in Mentalese, which somehow underpins human understanding of natural language conditionals. But the proofs we construct and check in logic and mathematics are not in Mentalese; they are in a formal language, or a mixture of natural language and mathematical notation. Thus the Distinction's friends are likely to resist identifying *a priori* knowledge with innate knowledge and *a posteriori* knowledge with acquired knowledge.

An illusion of depth

If the Distinction is shallow, why should it seem deep? A cynical answer is that its apparent depth is an artefact of philosophical conservativism: the Distinction draws prestige from its role in the epistemological tradition, going back to Kant and beyond; those who hold it in respect naturally regard it as deep. But even if there is something to such cynicism, it is not fully convincing. As a matter of common experience, sense perception is not part of every thought process. That is no mere quirk of human psychology; for computers too, internal computation is standardly distinguished from inputting and outputting. These elementary considerations may suggest that there *must* in principle be a radical distinction between knowledge dependent on sense perception and knowledge independent of it.

Such a distinction may indeed be drawn, if one looks only at the final cognitive stage culminating in the formation of the knowledge or belief at issue. But a taxonomy based just on the final stage would be blatantly superficial. For example, if one sees a black swan and thereby comes to know that not all swans are white, the knowledge gets classified as by thought rather than by perception, simply because the final stage was a deductive inference. No friend of the Distinction wants that result. Thus the taxonomy must take into account more distant epistemic progenitors too. That is straightforward enough when they are simply memories preserving relevant features of sense perception. But sense perception can also play a less overt role, sometimes by influencing which inferences one is disposed to make. For those who like the obscure term ‘intuition’, sense perception can influence which intuitions one is disposed to have. Analogously, inputs can modify a computer’s programme. In that case, there need be no input-independent outputs.

Does this amount to some form of *empiricism*, on which all knowledge and justification is *a posteriori*, perhaps on the lines sketched by Quine (1951)? Not at all. For empiricism characteristically minimizes the role of innate structure. If it cannot quite get away with positing an initial *tabula rasa*, it comes as close to that as it can, perhaps with an all-purpose Quinean similarity space for general inductive learning. By contrast, the present picture is consistent with richly differentiated innate structure, such as a Chomsky-style language module. Output may be highly sensitive to *both* inputs *and* a fixed cognitive architecture. Given such double sensitivity, the usual stereotype of the *a posteriori* is just as inept as the usual stereotype of the *a priori*.

The Distinction works well enough in an elementary introduction to epistemology, for a first-pass survey of the material to be understood. If we wish to go much deeper, the Distinction becomes an impediment to understanding.[3]

Notes

1 For internalist epistemologists, the primary distinction is between *a priori* and *a posteriori justification*. The difference is not crucial for the present arguments.

2 The inseparability of cognitive significance from representational form is arguably general, where even two synonyms differ in representational form (Williamson 2022c). We can track it by ascribing attitudes under representational guises, such as sentences in contexts. Thus someone may know the same proposition *a priori* under the guise 'Furze is furze' and *a posteriori* under the guise 'Furze is gorse'. For brevity and readability, I omit such qualifications. But an adequate account of a priori knowledge must reconcile it with the cognitive significance of the representational forms under whose guise we have such knowledge.

3 I was first alerted to problems with the Distinction through reflection on the epistemology of counterfactuals (Williamson 2007: 165-9). I developed the argument further in Williamson 2013. See Boghossian and Williamson 2020 for an extended debate, and Casullo 2022, Williamson 2022a and Melis and Wright 2022, Williamson 2022b for further exchanges. Many thanks to Joel David Hamkins, Daniel Kodsi, Jennifer Nagel, and audiences at the universities of Edinburgh, Lisbon, Rijeka, St Andrews, Texas Tech, and Turin (where it was given as the 2021 LLC Lecture) and the inaugural meeting of the British Society for the Theory of Knowledge at the 2022 Pacific Division meeting of the APA for useful questions and comments on earlier versions.

References

Boghossian, P., and Williamson, T. 2020. *Debating the A Priori*. Oxford: Oxford University Press.

Carruthers, P. 2015: *The Centered Mind: What the Science of Working Memory Shows Us About the Nature of Human Thought*. Oxford: Oxford University Press.

Castelvecchi, D. 2021: ‘Mathematicians welcome computer-assisted proof in ‘grand unification’ theory’, *Nature*, 595: 18-19.

Casullo, A. 2022: ‘A defense of the significance of the a priori–a posteriori distinction’, in Dodd and Zardini 2022.

Casullo, A., and Thurow, J. (eds.) 2013: *The A Priori in Philosophy*. Oxford: Oxford University Press.

De Toffoli, S. 2017: ‘“Chasing the diagram”: the use of visualization in algebraic reasoning’, *Review of Symbolic Logic*, 10: 158-186.

Dodd, D., and Zardini, E. (eds.) 2022: *Beyond Sense? New Essays on the Significance, Grounds, and Extent of the A Priori*. Oxford: Oxford University Press.

Fodor, J. 1975: *The Language of Thought*. New York: Thomas Crowell.

Griffiths, P. 2020: ‘The Distinction Between Innate and Acquired Characteristics’, *The Stanford Encyclopedia of Philosophy* (Spring 2020 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2020/entries/innate-acquired/>.

Kawai, N. 2019: *The Fear of Snakes: Evolutionary and Psychobiological Perspectives on our Innate Fear*. New York: Springer.

Melis, G., and Wright, C. 2022: ‘Oxonian scepticism about the a priori’, in Dodd and Zardini 2022.

Quine, W. 1951: ‘Two dogmas of empiricism’, *Philosophical Review*, 60: 20-43.

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 and Thurow 2013: 291-312. Reprinted in Boghossian and Williamson 2020: 117-136.

Williamson, T. 2022a. ‘Reply to Casullo’s defence of the significance of the a priori–a posteriori distinction, in Dodd and Zardini 2022.

Williamson, T. 2022b: ‘More Oxonian scepticism about the a priori’, in Dodd and Zardini 2022.

Williamson, T. 2022c. ‘Metametaphysics and semantics’. *Metaphilosophy*, forthcoming.