

FOURTH EXCURSUS

Warrants and Support Structures

So far, I have often cast notions such as apriority, scrutability, and knowability in modal terms: that is, in terms of what it is possible to know or to know a priori. We have seen that modal idealizations can lead to difficulties in some circumstances: they cause problems for propositional apriority and scrutability in cases of semantic fragility (see the end of the third excursus), and they cause other problems if there are brute modal constraints on possible reasoners (see section 2.7). It is also arguable that modal idealizations are not explanatorily fundamental: even when a scrutability thesis involving a modal idealization is true, it derives from more fundamental epistemological facts. So it is worth exploring nonmodal ways of understanding these notions.

In the introduction, scrutability theses are cast in terms of what one is in a position to know. This notion can be cashed out in modal terms, but it can also be cashed out in other ways.

In particular, a relevant notion of one's being in a position to know p can be cashed out in terms of there being a *warrant* for one to believe p . A warrant is a knowledge-apt justification, or a justification suitable for knowledge. There can be a warrant for one to believe p even if one does not in fact know or believe p . For example, when there exists a proof of p , this yields a warrant for believing p regardless of whether anyone proves p . These warrants are a form of *propositional justification*: a justification that supports belief in p for a subject, whether or not the subject believes p . This notion is standardly distinguished from *doxastic justification*: justification on which someone's justified belief in p is based.

On one notion of propositional warrant, one says that a subject *has* a propositional warrant to believe p when the warrant is (in some sense) within the subject's grasp. In this sense, the mere existence of a complex proof for p does not entail that a mathematically ignorant subject has a warrant for believing p . For our (idealized) purposes, however, this notion is too strong. The more relevant notion for our purposes is that of there *being* a propositional warrant for a subject to believe p (or more briefly, there being a propositional warrant for p), whether or not the subject has that warrant. This notion does not come with the requirement that the warrant is within the subject's grasp. There can be a warrant

for one to believe p even when knowing or believing p is beyond one's cognitive capacities. For example, even when a proof of p is enormously complex, it yields a warrant for a mathematically ignorant subject to believe p . We might call the first sort of warrant a *non-ideal warrant*, and the second sort an *ideal warrant*.¹

The framework of warrants makes a difference in cases of semantic fragility.² When p is the proposition expressed by the semantically fragile sentence S discussed at the end of the third excursus ('Snow is white iff actually snow is white'), one can argue that there exists a proof of p even though it is impossible to use it to prove p . In particular, there exists an abstract proof of S using the logic of 'actually'. S expresses p in the actual world, so this abstract proof of S is also an abstract proof of p . But if one were to *use* the proof to prove S , S would express p' rather than p , so one would not prove p .

What goes for proof goes also for warrant. In this case, there exists an (ideal a priori) warrant for believing p even though the warrant cannot be used to know p . This warrant is a propositional warrant that cannot be used as a doxastic warrant. If the subject is sophisticated and the proof is easy, the subject may even *have* a (non-ideal a priori) warrant for believing p , even though the subject cannot use that warrant to believe p . This last issue depends on precisely how one understands the conditions for having a (non-ideal) warrant, but either way, there will certainly *be* an a priori warrant for believing p . So there can be an a priori warrant for a subject to believe p even when it is not possible to know p a priori.

This suggests a nonmodal conception of propositional apriority: p is a priori in the nonmodal sense when there is an a priori warrant for some subject to believe p . Propositions expressed by semantically fragile sentences such as ' S iff actually S ' above may be a priori in this sense even if they are not a priori in the modal sense. Likewise, this suggests a nonmodal conception of knowability: p is knowable by a subject in the nonmodal sense if that subject has warrant for believing p .

One can define nonmodal apriority and knowability for sentences in a similar way: S is nonmodally a priori if there is an a priori warrant for some subject to believe S (where believing S is understood as in the third excursus), while S is nonmodally knowable for a subject if there is a warrant for that subject to believe S . Unlike the propositional versions, these nonmodal notions of sentential

¹ One might require that even an ideal warrant is within a subject's idealized grasp (in some nonmodal sense), where the idealization involves idealization of reasoning. So non-ideal and ideal warrants could both be seen as providing 'subjective' reasons (reasons for a subject that are in some sense available to the subject), although the latter is a highly idealized variety. Thanks to John Bengson, Jon Kvanvig, Nico Silins, and Chris Tucker for discussion of warrant.

² The next few paragraphs will make most sense in light of the last few paragraphs of the previous excursus, on semantic fragility (or alternatively, in light of 'Actuality and Knowability'). Readers can either read that material first or alternatively skip directly to 'What is a warrant, exactly?' below.

apriority and knowability will not come apart from modal apriority and knowability in cases of semantic fragility. They may still come apart if there are brute modal constraints on possible thinkers, however. For example, if no possible thinker can carry out a proof that involves more than a million steps, then certain true mathematical sentences (and propositions) will be a priori in the non-modal sense but not in the modal sense.

One can also appeal to warrant to define nonmodal notions of scrutability. For example: q is inferentially scrutable from p when knowing p would provide a warrant for q ; q is conditionally scrutable from p when one has a warrant for accepting *if p , then q* ; and q is a priori scrutable from p when there is an a priori warrant for accepting *if p , then q* . These definitions are unaffected by semantic fragility. In the cases from the third excursus on which a semantically fragile sentence P is scrutable from Q on a modal definition of scrutability, the corresponding proposition p is not scrutable from q on a modal definition of propositional scrutability, but p will plausibly be scrutable from q on the nonmodal definition of propositional scrutability.

I often cast notions such as apriority and scrutability in modal terms elsewhere in this book, in part because modal analyses of notions such as apriority are more familiar than those in terms of warrant, and partly because the problem of semantic fragility does not affect sentential scrutability. But it is reasonable to hold that the warrant-based notions are more fundamental than the corresponding modal notions. When it is possible to know p a priori, this is typically possible because there is an a priori warrant for p . Likewise, when S is scrutable from C in a modal sense, this is typically because S is scrutable from C in a warrant-based sense.

What is a warrant, exactly? A warrant is plausibly a sort of justification: a justification suitable for knowledge. But what is a justification? On some views, there is no particular entity that is a justification, but simply a relation between subjects and propositions misleadingly labeled ‘There is a justification for s to believe p .’ For our purposes, however, it is useful to develop a more substantive understanding.

I will develop an understanding of warrants and justifications as *support structures*. This understanding is inspired by the special case of proof. When belief in p is warranted by a proof, the corresponding warrant for p derives from the structure of propositions in the proof, with support relations corresponding to logical steps. We can count this structure as a support structure. More generally, a support structure will involve directed graphs of labeled propositions with support relations between them, capturing the justification for a proposition.

It is more common in the epistemological literature to view warrants and justifications as propositions or perhaps as sets of propositions. But these propositions only play their role in virtue of their position in a support structure, and for various purposes it is useful to make this structure explicit. The case of a

proof suggests that there can sometimes be a warrant (in an intuitive sense) for someone to believe a proposition even though there is no clear proposition or set of propositions that constitute the warrant. It does not seem right to say that the warrant is the proposition expressed by the first step of the proof or that it is the proposition expressed by the penultimate step. It is also not perspicuous to say that it is the set of all propositions in the proof thrown together, as so many of these propositions are themselves supported by other propositions in the set. Viewing warrants as support structures avoids these problems. This view need not involve a substantive disagreement with those who view warrants as propositions: to be fully explicit we could call these structures ‘warrant structures’ instead of warrants, and likewise for justifications. The details of this framework are not crucial for addressing the problem of semantic fragility, but they will play a role for other epistemological purposes in chapters 3 and 4. So I will develop the framework in some detail.

We can start with doxastic justification. Take any justified belief that p . Something (a belief or an experience, for example) *justifies* belief in p when it supports p or provides evidence for p in a way that yields prima facie justification for the belief that p . It *directly justifies* belief in p when it justifies belief in p and does not justify belief in p wholly in virtue of justifying something else that justifies belief in p .³

A *direct justification* for p can be represented as a graph consisting of a node for p along with nodes for elements (if any) that directly justify belief in p , with arrows from those nodes to p . For example, when a belief in p is justified inferentially, it will be directly justified by inference from one or more other justified beliefs: a belief that q , a belief that r , and so on. Here we can say that these other beliefs, collectively, provide a direct justification for the belief that p . A direct justification can here be represented as a node for p with nodes for q , r , and so on, with arrows from them to p . If the justification is redundant, so that the belief that q suffices on its own to inferentially justify the belief that p , there will also be a direct justification consisting of a node for q with an arrow to a node for p .

When a belief in p is justified non-inferentially, either it will be directly justified by some evidence distinct from the belief (e.g. perceptual evidence) or it will be justified by no such evidence (as on some views of basic belief). In the first

³ The justification relation is an epistemic grounding relation, and should be distinguished from the metaphysical and conceptual grounding relations discussed in the excursus on grounding and elsewhere. For something to stand in this relation to p is not for it to be the metaphysical grounds for the belief that p or the metaphysical grounds for the justification of the belief that p . For example, on a reliabilist view, a reliable process might serve as metaphysical grounds for the justification of a basic belief that p , but this process will not itself stand in the epistemic grounding relation to the belief that p . When a belief that p directly justifies a belief that q , however, it is plausible that their standing in this support relation (or the metaphysical grounds of their standing in this relation) will serve as part of the metaphysical grounds for the justification of the belief that p .

case, we can say that p is justified non-inferentially and evidentially. Here a direct justification for the belief that p can be represented as a node for p with arrows from nodes for any directly justifying evidence. In the second case, we can say that the belief that p is justified non-inferentially and non-evidentially. Here a direct justification for p can be represented by a node for p alone, with p labeled as a basic belief. If the belief that p is self-justifying (if this is possible), the graph will include an arrow from p to itself. If there are beliefs that are justified both inferentially and non-inferentially, then both sorts of support can be included in the structure; though as above, in cases where the justification is redundant, there will also be direct justifications that exclude redundant elements.

An *indirect* justification for p will include a direct justification for p and also justifications for one or more elements in the structure that supports p . A *justification* for p is a direct or indirect justification for p . A *full* justification is a justification that includes a justification for every belief in the structure. A *partial* justification is a justification that is not a full justification.⁴

For example, when belief in p is inferentially justified by belief in q , and belief in q is inferentially justified by belief in r , and belief in r is non-inferentially justified by evidence e , a direct justification will include only a link from q to p . A full justification p will include links from e to r to q to p . There will also be an indirect partial justification with links from r to q to p .

If there can be circles of justificational support (e.g. from p to q to r to p), then justifications can include these circles. If there can be infinite chains of support, then justifications can include these chains. When there are no such circles or infinite chains, we can say that a justification is *classical*. The *grounds* of a classical justification are its initial elements. When a belief that p has a classical justification (whether partial or full), we can say that the grounds of that justification *ground* belief in p . When a belief that p has a full classical justification, we can say that the grounds for the justification *fully ground* belief in p . The grounds of a full classical justification will be *basic evidence*: these may include basic beliefs and/or basic nondoxastic evidence.⁵

For simplicity, I will adopt a model on which evidence always involves propositions. So introspective evidence might involve the proposition that one is in a

⁴ Note that a partial justification is partial in the sense that it omits some elements that play a justifying role, but not in the sense that it yields support that merely weighs in favor of p without justifying belief in p . I am taking it that all of the justifications I consider here are strong enough to justify belief in p . The elements omitted in a partial justification will either be redundant elements (whose contribution is not required for justification) or indirectly justifying elements (whose contribution is mediated by another element).

⁵ Even when there is no full classical justification for p , there may still be a partial classical justification for p . For example, when full justification involves an infinite chain, there will always be a partial justification without such a chain. Even when full justification involves a circle, it may be that there is a partial justification without a circle.

given mental state: for example, the introspective evidence for the belief that one is in pain may be the proposition that one is in pain. Likewise, perceptual evidence may involve the propositional contents of perceptual experience: for example, the perceptual evidence for the perceptual belief that there is something red in front of one may be the (perceptually represented) proposition that there is something red in front of one.⁶ Something similar may go for evidence provided by (for example) intuition, memory, and testimony, if one thinks that these involve sources of non-inferential justification distinct from perceptual or introspective evidence.

To distinguish the different roles for various propositions here, the structure will label a proposition p as a belief proposition (p is believed), perceptual evidence (p is the content of a perceptual state), introspective evidence (p is introspectively experienced, or perhaps need only be true), and so on. So in the case above, the justification might look like: e (experienced) $\rightarrow p$ (believed) $\rightarrow q$ (believed) $\rightarrow r$ (believed). This can be read as saying that experience as of e justifies belief in p , which justifies belief in q , which justifies belief in r . The model can easily be adapted to views on which evidence and the relata of the support relation are mental states rather than propositions. On such a view the nodes in the structure will be mental states such as beliefs that p or even nonpropositional mental states such as pains. The model can also be adapted to views on which some of the relata are facts (or perhaps property instances or other entities) rather than propositions.

All this is a sketchy and incomplete picture of support structures. For completeness, the picture should be elaborated in various ways. Importantly, one can also allow support relations themselves to be supported or defeated by evidence. As I discuss in the section on empirical inference in chapter 4, an inferential relation between s and t (say, a direct inference from someone's being a bachelor to his being untidy) can itself be grounded in prior perceptual evidence e . This can be represented by an arrow from e to the arrow between s and t . Similarly, as discussed in chapter 3, an empirical recognitional capacity connecting a perceptual experience e to a recognitional belief b (e.g. recognizing a certain shape as an iPhone) may itself be grounded in prior evidence e' (say, evidence that iPhones have that shape). As I also discuss there, it is even possible to allow that perceptual experiences can sometimes be supported by prior evidence (so that seeing a person as angry may be grounded in prior evidence that they look that way when angry). If so, experiences can sometimes occupy non-initial positions in support

⁶ I am using 'evidence' in a broad sense on which all epistemic grounds count as evidence, as opposed to a sense on which only justifiers that are true or that are known count as evidence. Whatever one says about 'evidence', I think it is plausible that false propositions can serve as direct justifiers: for example, a belief that p or an experience as of p can justify a belief that q whether p is true or false.

structures. Adding structure to support relations is also possible, for example to indicate that p and q conjointly support $(p \& q) \vee r$ while r supports it separately.

A propositional justification is the same sort of item as a doxastic justification, with the difference that there can be a propositional justification for a subject to believe p even without the subject believing the belief propositions in the justification. It is still plausibly required that for any perceptual and introspective evidence specified in the support structure, the subject must have (or at least have had) corresponding mental states. (If we understand the nodes of the support structure as mental states, the evidential nodes will be mental state tokens while the supported nodes will be mental state types.) A propositional justification yields a doxastic justification when the subject has all the relevant beliefs, and when the beliefs are properly based on each other and on the evidence in a way that reflects the support relations in the structure.

A (propositional or doxastic) warrant is a (propositional or doxastic) justification that meets certain further conditions to make it knowledge-apt for the subject. These arguably include the conditions that all the included propositions be true, that any initial belief propositions (especially in a doxastic warrant) be known, and that there are no defeaters and no Gettier circumstances. These conditions might be varied or extended, but I will stay neutral on the precise conditions required.⁷

The model has a foundationalist flavor to it, but it does not presuppose foundationalism. As long as a view acknowledges the distinction between inferential and non-inferential justification (even if it holds that all justification falls on one side), the model will be coherent. Coherentist and infinitist views will allow nonclassical justifications, and may or may not give a role to basic evidence. Reliabilist and other externalist views may sometimes find non-inferential

⁷ What is the relation between these notions and the standard notions of subjective and objective reasons for belief? Let us say that a *reason structure* is a justification freed of the requirement that the subject have states corresponding to the relevant perceptual and introspective evidence. Then a reason structure provides subjective reason to believe p roughly when the subject either has mental states corresponding to the initial elements in the structure or perhaps when the subject is in a position to have them (at least where these elements correspond to a priori beliefs). A reason structure provides objective reason to believe p when all the initial elements are true. All of the justifications I have talked about yield at least subjective reasons (although these may be idealized subjective reasons in some cases). Warrants yield both subjective and objective reasons. Here reason structures are roughly analogous to valid arguments, subjective reasons are roughly analogous to valid arguments whose premises are justifiably believed, objective reasons are roughly analogous to sound arguments (valid arguments whose premises are true), and warrants are roughly analogous to sound arguments whose premises are known. (The fact that experiences need not be objects of justification complicates the analogy, however, as does the fact that not every argument transmits justification.) One might also develop a more general notion of a ‘basing structure’, analogous to an argument (whether valid or invalid): such a structure might reflect only the basing (or potential basing) of certain beliefs on others, whether or not this basing goes along with justification.

justification where other views find inferential justification, and may or may not give a role to non-inferential evidential justification. (If a view does not recognize the notion of non-inferential evidential justification, we can count it as classifying all non-inferential justification as non-evidential.) Speaking for myself, I think it is enormously plausible that there is much inferential justification and much non-inferential evidential justification, so I think that full justifications will often be quite complex.

We can use this framework to help analyze the distinction between a priori and a posteriori justification. At least among full classical justifications, an a posteriori justification will be one with some empirical grounds, while an a priori justification will be one with no empirical grounds. Empirical grounds will include all perceptual and introspective evidence, and perhaps other basic evidence depending on one's views. If one accepts (as I do) that all basic empirical evidence is perceptual or introspective evidence, we can say more simply that a justification is a posteriori iff its grounds include perceptual or introspective evidence. If one holds that there are other sorts of basic empirical evidence, one will need further criteria for classifying basic evidence as empirical or non-empirical. If there are full nonclassical justifications, one will need further criteria to classify these as a priori or a posteriori; here the existence of an empirical ground will serve at least as a sufficient condition for such a justification to be a posteriori.⁸

To analyze a priori doxastic justification, we can then say that a subject's belief is justified a priori if it has an a priori doxastic justification. To analyze a priori propositional justification, we need only invoke the idea that there is an a priori justification for a subject to believe a proposition. We can make parallel claims about a posteriori justification, and about a priori and a posteriori warrant. We can distinguish a special class of conclusive a priori and conclusive a posteriori justifications—that is, justifications for certainty rather than for mere belief—by requiring conclusive basic evidence and conclusive support relations in a justification.

We can also distinguish some a priori warrants as *conceptual* warrants, where a conceptual warrant for p is a conclusive a priori justification that derives from the concepts involved in p . I develop one way of understanding this notion, on an inferentialist model of concepts, in the seventeenth excursus. We might then

⁸ When support relations are themselves empirically supported, as with empirical inferences, the relevant empirical evidence will count as grounds and the resulting beliefs will be a posteriori. There are tricky cases discussed at the end of E8 involving processes (perhaps innate processes) that are not obviously grounded in experience but that are not traditionally a priori either. A further option to handle these cases is to label some support relations as a priori relations (perhaps based on the positive character of the support) and to require that an a priori justification involve only a priori support relations.

see an analytic truth as one that expresses a proposition for which there is a conceptual warrant.

We can straightforwardly extend the current analysis to warrant for sentences and for thoughts. One can also extend the current model to analyze cases in which prima facie justification is defeated or in which support is only partial. Finally, one can extend it to analyze justification for having certain credences in propositions (or sentences or thoughts), based on credences in other propositions (or sentences or thoughts) and evidence. The last analysis can yield a warrant-based analysis of rational credences, helping to avoid the problems for modal analyses discussed in the previous excursus.