

6

Loosening the Need-Concept Tie

Edward Craig

What drives people to develop the concept of knowledge? There seems to be a word for 'know' in every language, which suggests that it answers to highly general needs of human life.¹ But what are those needs? This is the question that leads Edward Craig (1990) to offer a genealogy of the concept of knowledge. He takes his cue from Williams (1973a, 146), who pointed out that philosophers have traditionally suffered from a *déformation professionnelle* in thinking about knowledge: they have considered it from the perspective of the *examiner* who assesses whether someone knows something already known to the examiner ('Does A know that *p*?'). Yet as Williams points out, the examiner's situation is a rather peculiar one. The more natural home for the concept of knowledge is the situation of the *inquirer* who seeks to identify someone who knows what the inquirer does not yet know ('Who knows whether *p*?'). Modelling the inquirer's situation in the state of nature, Craig's genealogy describes how the need for true beliefs about the immediate environment would lead the inquirer to develop a primitive, prototypical form of the concept of knowledge as a tool by which to identify what, given the individual's needs and capacities, would be good informants. It then describes how the concept gradually loses this indexing to subjective needs and capacities to become more like our concept of knowledge.

Craig's pragmatic genealogy is the first book-length investigation of its kind that is entirely devoted to a single concept, and it is carefully focused on the *concept* of knowledge rather than on knowledge itself. *Pace* Michael Hannon, who thinks 'it makes little sense to distinguish the attempt to become clearer about our concept of knowledge from the attempt to become clearer about knowledge as such' (2019, 32), this distinction is crucial to the method of pragmatic genealogy. A pragmatic genealogy of knowledge itself would be quite a different affair, leading us to ask why a creature would need to *have* knowledge *about* its environment rather than why it would need to become sensitive to the *presence of* knowledge *in* that environment.

Perhaps counterintuitively, a pragmatic genealogy of the concept of knowledge can remain instructive even if we take knowledge itself as given. We can

¹ See Wierzbicka (2018).

simultaneously illustrate this important methodological point and render Craig's genealogical approach more attractive by examining its alleged incompatibility with the currently more influential approach of 'knowledge-first' epistemology (KFE) propounded by Timothy Williamson and others. KFE involves two claims: (a) knowledge is the most general factive stative attitude,² an attitude one necessarily bears to a truth if one bears any other factive stative attitude to it; and (b) the concept of knowledge is an unanalysable primitive that should act as a basis for explaining related concepts like *believing*, *seeing* and *remembering*.³

There are three reasons why one might think Craig's genealogy incompatible with KFE.⁴ The first is pointed out by Williamson himself. Though he applauds Craig's dismissal of the traditional programme of analysing the concept of knowledge, Williamson still considers Craig's project marred by its failure to acknowledge that the need for knowledge is prior to the need for true beliefs:

[Craig's project] remains too close to the traditional programme, for it takes as its starting point our need for true beliefs about our environment . . . , as though this were somehow more basic than our need for knowledge of our environment. It is no reply that believing truly is as useful as knowing, for it is agreed that the starting point should be more specific than 'useful mental state'; why should it be specific in the manner of 'believing truly' rather than in that of 'knowing'?

(Williamson 2000, 31n3)

Call this the *wrong-starting-point* problem.

The second reason is that Craig's original prototype of the concept of knowledge—*proto-knowledge*, as we might call it following Martin Kusch (2009b)—is a concept tracking 'proto-knowers' or good informants, and as Miranda Fricker highlights, Craig 'tends to describe the good informant as someone recognizable as having a true belief' (2007, 144n17). But if a proto-knower is someone who, among other things, must have a true belief that *p*, this puts the concept of belief at the heart of the prototype of the concept of knowledge; and if this prototype is still part of our concept of knowledge, as Craig claims (2007, 191), this conflicts with one of the central tenets of KFE—that the concept of knowledge does precisely *not* involve the concept of belief. The problem, then, is that if Craig defines the good informant as someone who truly believes that *p*, this ultimately puts the concept of belief at the core of the concept of knowledge, 'and therefore depicts belief as prior

² An attitude being *factive* if it is one that can only be borne to a truth, and *stative* if that attitude is a state.

³ Williamson (2000, 33, 44, 185).

⁴ See also Hannon (2019, 20–2) for a comparison of knowledge-first epistemology and Craigean function-first epistemology which emphasizes the advantages of the latter over the former.

to knowledge, so that knowledge is conceived as true belief plus a bit' (Fricker 2007, 144n17).⁵ Call this the *believing-as-the-core-of-proto-knowing* problem.

Third, and perhaps most importantly, KFE insists that the concept of knowledge is primitive and explanatorily fundamental (Williamson 2000, 185); but Craig's genealogy seeks to explain the concept of knowledge in terms of something more primitive—in particular, subjective needs for a concept that is more primitive than the concept of knowledge. This also seems incompatible with KFE, since, if the concept of knowledge is indeed fundamental in the way Williamson proposes, 'there is little space for any genealogy' (Kusch 2011, 12).⁶ Call this the *no-room-for-genealogy* problem.

My first aim in this chapter is to show, against these three objections, that Craig's genealogy is compatible with KFE and even lends limited succour to it. Because Craig undertakes to explain the *prima facie* plausibility of extant conceptual analyses of knowledge, he remains overly beholden to the terms in which they are given—in particular, to their preoccupation with true belief. But if Craig's genealogy is freed from the trappings of past proposals, it can dispense with the notion of true belief, reach its conclusions all the same, and even lend succour to KFE. In showing how this can be so, I vindicate an otherwise enigmatic suggestion made by Williams. Williams took Craig's book to *support* KFE, because what it really indicated, on Williams's view, was that knowledge is prior to belief.⁷

My second aim is to bring out two respects in which, beyond what it tells us about the concept of knowledge, Craig's genealogy empowers the methodological framework it instantiates by identifying practical pressures driving what I call the *de-instrumentalization* of concepts—the process whereby concepts are driven to shed the traces of their practical origins. Craig's distinctive contribution to the pragmatic genealogical method is thus to loosen the need-concept tie: he demonstrates that far from reducing our mental life to instrumental thought, pragmatic genealogy not only accommodates but in fact predicts the de-instrumentalization of concepts.

My third aim, finally, is to show how Craig harnesses pragmatic genealogy's ability to help us assess and synthesize competing accounts of concepts. In coming to understand what concepts would serve us well, we gain a sense of what our concepts are likely to be and how competing accounts of them relate to each other. As we shall see, this provides a valuable way to resolve philosophical stand-offs.

⁵ Fricker (2007, 144n17) goes on to remark that it would in principle be open to Craig to avoid this problem. I argue for the stronger claim that this problem in fact does not arise in the first place.

⁶ Kusch's own position is that this problem can be overcome: see Kusch (2009b, 90) and Kusch and McKenna (2020, 1068–70).

⁷ For the claim that Williams took Craig's book to indicate that knowledge is prior to belief, see Fricker (2007, 113–14n9). Fricker herself endorses a reading on which Craig can in principle remain agnostic on the priority question (2007, 144n17).

6.1 Fictional Starting Points

Let us begin with the *wrong-starting-point* problem: Craig starts from the thought that ‘human beings need true beliefs’ (1990, 11) rather than from the thought that human beings need knowledge. Does this vitiate the entire project from the perspective of KFE? I think not, and elaborating why not will take us through several important methodological considerations bearing on the choice of starting points for pragmatic genealogies.

The essential first move in Craig’s genealogy is not its focus on the need for true beliefs, but its focus on the practical problem one faces in the situation of the inquirer. You will be in the situation of the inquirer if you are ignorant as to whether p and need to find out, if you are to be successful in navigating your environment and satisfying your most basic needs, whether p —whether, for example, the bear went into the cave. This means that we can in principle preserve neutrality between belief-first and knowledge-first accounts without prejudice to Craig’s account: we can say that human beings need truths about their environment, leaving it open whether the attitude they must bear to these truths is one of believing or knowing. Inquirers simply face a problem of the form: how to come by the truth as to whether p ?

Arguably, however, careful scrutiny of the inquirer’s situation and its practical challenges already leads one to the conclusion that what the inquirer must seek is really knowledge rather than true belief. Williams gives an argument to this effect (whose form and spirit foreshadows Craig’s pragmatic genealogy) in his 1978 book on Descartes (2005b, 23–32). The argument turns on taking seriously the idea that the inquirer’s situation is a practical situation: instead of considering inquirers as disembodied intelligences and asking, at an entirely abstract level, what their respective attitudes to truths, beliefs, or knowledge should be, we consider what form the search for truth takes in a concrete practical context. For human inquirers, this context will notably be marked by the contingent but no less important fact that they operate in a non-magical world where true beliefs do not just pop up when they wish for them, but have to be sought out and identified as such. This imposes important constraints on what the search for truth can look like in practice. *Even if* we assume that our inquirers initially want truth in the form of true beliefs rather than knowledge, these constraints entail that seeking truths and seeking knowledge cannot ultimately be two separate enterprises, because once we consider how the search for true beliefs would have to be *operationalized*, we realize that in our world, that search must take the form of a search for knowledge.

More precisely, Williams’s argument runs as follows: to believe that p is already to take p to be true, so that inquirers cannot first seek to believe things and then decide which beliefs to treat as true; inquirers interested in arriving at true beliefs must *seek* true beliefs from the first. But in a non-magical world, inquirers do not just come

by true beliefs by wanting them; they need to *operationalize* the search for truth: to develop practically workable *methods* by which to track down true beliefs. However, truth is not the kind of property that inquirers might track directly; inquirers therefore need to track some *indicator property* which recognizably indicates, if only defeasibly, that the beliefs bearing that property are true. That indicator property, Williams writes, will have to be something like ‘*appropriately* produced in a way such that beliefs produced in that way are generally true’ (2005b, 31), where ‘*appropriately*’ is meant to rule out accidentally true beliefs, which, again from a practical point of view, inquirers cannot reasonably make it part of their *policy* to seek out, since the thing about accidentally true beliefs is that they might well have been false. Assuming that possession of that indicator property is sufficient to turn a belief into knowledge—something that KFE could accept as long as the transformation does more than just *add* something to belief—we reach the conclusion that what situated inquirers really need to seek, and what their concepts should equip them to seek, is nothing short of knowledge. Williams concludes that ‘starting merely from the idea of pursuing truth in a non-magical world, and so of the truth-seeker’s using methods of enquiry, we do arrive at the conclusion that the search for truth is the search for knowledge’ (2005b, 31). ‘In wanting the truth’, the inquirer already ‘wants to know’ (2005b, 27). This line of argument is one way to hear Williams’s remark that knowledge emerges as being prior to belief on the Craigean approach. The need for truths, when lowered into the practical context of human life where it is confronted with contingent hurdles that remain invisible from a purely abstract perspective, assumes the form of a need for knowledge.

But if our starting point is spelled out in these KFE-friendly terms, it arguably presupposes the existence of knowledge itself, and one might worry that this conflicts with the declaredly pragmatist spirit of Craig’s account⁸—in particular, that it runs afoul of the pragmatist commitment to eschewing, wherever possible, the metaphysics-inviting strategy of explaining the *concept* of X in terms of the *prior existence* of X *itself*. Must pragmatists worthy of the name not put a ban on mentioning the object of the concept they seek to understand, and try to explain how the concept would come about nonetheless, driven by forces that have nothing to do with being sensitive to its object?

The simplest response would be for the pragmatist to concede that knowledge is just one of those cases where the pragmatist ambition to explain X in terms of the concept of X without drawing on X in any way reaches its limits, and that this is indeed no surprise if, as advocates of KFE suggest, knowledge is fundamental to having a mind at all (Williamson 2000, 48). Within KFE, the idea of imagining a human community in which there is no knowledge ‘is probably not intelligible’,

⁸ Craig describes his method as ‘pragmatic synthesis’ or ‘practical explication’ (1990, 8, 141). He also speaks of his ‘pragmatic method’ (1993, 44).

because ‘human beings always have perception and other automatic ways of knowing that it would be hard to imagine us . . . lacking for any significant length of time’ (Reynolds 2017, 12).

But while it is true that *we* could not *grasp* what it would mean for a creature to have knowledge if we did not have the concept of knowledge—it would not be intelligible to us *as* knowledge—it does not follow that knowledge cannot *exist* without the concept of knowledge. Consequently, there is room for us to contemplate a situation in which agents have knowledge, but lack the concept thereof, so long as *we* possess the concept of knowledge that allows us to conceptualize the situation in this way; and we can illuminatingly do this as long as the practical contribution of the concept differs from the practical contribution of its object. This concessive response retains the compatibility of Craig’s genealogy with KFE while granting that its strategy is characteristic of pragmatism only up to a point.

But there are also two less concessive responses. The first insists that Craig’s genealogy does not run afoul of pragmatist principles at all: even on the KFE-friendly reading, Craig does not start by asking what knowledge *is*, or what ‘knowledge’ *refers* to; in good pragmatist fashion, he starts with agents and the practical needs that give point to the concept of knowledge. Moreover, as Simon Blackburn (2017) has shown, it is by no means clear what exactly the pragmatist ban on mentioning certain entities in one’s explanations should amount to. What is clear is that Craig *would* run afoul of pragmatist principles if he said that we think in terms of knowledge because there is so much knowledge around—if, to put it slightly more carefully, his explanation had the following form:

There is X.

We need to be suitably sensitive to the presence of X.

Therefore, we develop the concept of X for this purpose.

But Craig explicitly rejects a strategy along these lines (1990, 3), and this is where he proves himself a pragmatist after all. It is not the object—knowledge—that *attracts* the use of the concept of knowledge. It is the agent’s needs that *drive* the emergence of the concept by which the object, knowledge, can then be delineated. In other words, Craig’s genealogy is subtler and more informative than the above schema because the need to be sensitive to the presence of knowledge is his *explanandum* rather than his *explanans*.⁹ He derives this need from other needs instead of assuming it.

What makes model-based genealogies so useful in dealing with concepts that are fundamental to human life really only comes out once we take an even less

⁹ I leave aside the question whether knowledge’s being a social or a natural kind makes any difference to the compatibility of KFE with Craigean genealogy. I agree with Kusch (2013) and Kusch and McKenna (2020, 1065–70) that it does not.

concessive line, however. We can freely grant that if the concept of knowledge is as central to our lives as KFE claims, any state-of-nature scenario will soon covertly presuppose not just the existence of knowledge, but even the existence of the *concept* thereof: given KFE, any actual community that lacked the concept of knowledge would also have to lack a host of further concepts and practices that *depend* on the concept of knowledge, such as the practice of *telling* someone that *p* (governed by the norm that one must know that *p*), or even the concepts of *seeing* or *remembering* that *p* (which imply knowing that *p*). This is a point pressed by Elizabeth Fricker (2015) in her critique of Craig. In response, Steven Reynolds (2017) has suggested that we imagine a community deploying variants of the concepts of *seeing* or *remembering* that do not depend on the concept of knowledge—*seeing** and *remembering**, which can be used to report appearances without entailing attributions of knowledge (2017, 31–4). In principle, one could do this across the board until one had eradicated all traces of the concept of knowledge in the initial stage of one’s model. But surely what matters is not that we *do* replace all affected conceptual practices with unconnected variants in our model, but that we *could* do so without losing the genealogy’s explanatory force. Just because the genealogy is a model that idealizes and sometimes distorts reality, we can acknowledge that some things cannot, ultimately, be separated, and still ‘fruitfully postpone considering all these things together’ (2000, 153), as Williams puts it. We can instructively *separate in fiction what belongs together in reality*, thereby treating as independent what is in fact related, in order to identify what needs in particular a certain element in our conceptual arsenal answers to.

By allowing us to consider as arising separately what in reality has to arise together, pragmatic genealogy helps us organize and break down into manageable pieces the practical contribution made by concepts that are fundamental to human life. KFE claims for the concept of knowledge a status that many philosophers are willing to grant the concept of truth, namely the status of being so fundamental that we struggle to imagine language-using human societies without already covertly drawing on it. But why should the mere fact that the concept of truth is fundamental in this way bar us from using state-of-nature fictions to help us identify some of the functions that the concept of truth performs? As Miranda Fricker points out, it is precisely the ‘genius of using the state-of-nature format in the arena of epistemology’ that it allows one ‘to tell a narrative story about X (e.g. the concept “know”) even where we find it otherwise barely intelligible that there could have been a narrative development towards X’, for instance because ‘the idea of a progression towards X is conceptually impossible’ (1998, 165).¹⁰ The fact that a concept is so fundamental as to be involved in some of the most basic human activities only means that we would do well not to think that one simple state-of-

¹⁰ Fricker notes that Williams made this point in a lecture entitled ‘Truth and Truthfulness’, delivered at Birkbeck in May 1997 (Fricker 1998, 165n13).

nature story will exhaust that concept's practical contribution across the entire range of human affairs. The question, then, is not whether the concept that forms the target of the genealogical narrative was covertly presupposed in some part of the state-of-nature model—if the concept is fundamental enough, it very likely was—but whether it was presupposed *in a way that renders the narrative unilluminating*. As long as the genealogy uncovers instrumental relations between the concept and certain needs that we were unaware of, it will retain its interest.

Even by pragmatist lights, therefore, Craig can illuminatingly ask what would drive a community of inquirers who need to know whether *p* to develop the concept of knowledge—even if this concept is internally related to other basic human activities, and even if this presupposes the existence of knowledge. The former is no problem because the genealogy is imaginary, and can, in virtue of this, help us identify some of the salient ways in which the concept serves our needs; and the latter is no problem because the genealogy does not presuppose the need to be sensitive to the presence of knowledge, but rather, as we shall now see, illuminatingly derives this need from other needs.

6.2 What Informants Need to Be

Let us turn to the *believing-as-the-core-of-proto-knowing* problem. If the good informant must be someone who truly believes that *p*, this puts the concept of belief at the core of the concept of proto-knowledge, and thereby also at the core of the concept of knowledge. Upon closer inspection, I want to suggest, this problem does not arise, and Craig's account even lends succour to KFE.

To see why, we need to delve further into Craig's genealogy. Given inquirers' need to find out the truth as to whether *p*, they can make some headway by relying solely on perception, reasoning, and memory—their 'on-board' resources (Craig 1990, 11). Yet the mere fact that inquirers are not all in the same place at the same time entails that there are pressures on inquirers to engage in cooperation to find things out. Someone else may have seen what I have not. Hence, inquirers need to tap into one another's stores of information: where a given inquirer's on-board resources are no help, the inquirer needs a *good informant* as to whether *p*. In characterizing good informants, Craig also tends to take his lead from the programme of conceptual analysis. He notes that they typically are individuals who truly believe that *p* and also display further properties that render them suitable as informants (1990, 12–15, 96). But what determines which characteristics are central to his account is whether they are practically relevant to the inquirer's success. 'We must never forget', he insists in an echo of Williams, 'that the inquirer's situation is a practical one' (1990, 15). From a practical point of view, good informants need to be (i) as likely to be right as to whether *p* as is necessary for the inquirer's purposes; (ii) accessible to the inquirer here and now;

(iii) intelligible to the inquirer; (iv) honest, in the most basic sense of being open with the inquirer; and (v) such that, for whatever reason, the inquirer finds their testimony convincing.

If inquirers need good informants, they also need to be able to recognize them as such. Good informants need to be identifiable as satisfying some or all of the above conditions through indicator properties: standing in the right causal relation to the state of affairs in question, for instance, or having proven reliable in the past, or being able to offer justifications. It is here that we begin to discern the roots of the vexed question of warrant, i.e. of what the ‘third condition’ on knowledge should be. The genealogy sheds light on this third condition ‘by setting it against helpful ideas of what the point of imposing such a condition might be’ (Williams 2010, 215). As Craig puts it, we hypothesize the point of something like the concept of knowledge, combine this with ‘a few platitudes about the way in which human inquirers operate’ (1990, 16), and generate a set of descriptive conditions characterizing the circumstances under which the concept *typically* serves that point.

Given that inquirers need to become sensitive to the presence of certain indicator properties, they need a *concept* whose application tracks these indicator properties. It is here that the prototypical form of the concept of knowledge—*proto-knowledge*—comes in: it fills the need for a concept picking out ‘good informants whether *p* given the needs and capacities of the inquirer here and now’. Note that *proto-knowledge* is still markedly different from our concept of knowledge. It is strongly perspectival, i.e. indexed to the particular inquirer’s needs and capacities at that time and place: a proto-knower is someone who will suit *my* needs, given *my* capacities, *here* and *now*. Moreover, the concept is what we may call *purely instrumental*: a concept is purely instrumental just in case (a) the concept is instrumental in serving the concept-user’s needs; and (b) the concept is the concept *of* something instrumental in serving the concept-user’s needs. The status of being a proto-knower is so closely tied to my needs that there is no conceptual room for a proto-knower who, for whatever reason, is no use to me.

We can now see that the *believing-as-the-core-of-proto-knowing* problem evaporates under analysis. This is because what the inquirer has a need *for*, on Craig’s account, is not in the first instance someone who has a *true belief* as to whether *p*, but someone who *proto-knows* whether *p*, and believing that *p* is not a necessary condition on proto-knowing that *p*. It is only one among many properties that correlate well with telling the truth about *p* (1990, 13). Circumstances concurring, good informants may not need to believe what they say, and even a diffidently offered piece of information can come to be accepted by the inquirer—students might come to know the theory of evolution through their competent but creationist schoolteacher, for example (Lackey 2007). It is simply more likely that an informant who believes that *p* will come out with *p* and do so in a manner capable of persuading the inquirer. But from a practical point of view, what matters is that

the informant is likely enough to be right, accessible, intelligible, and willing to say whether p in a manner capable of persuading the inquirer.

The heart of Craig's proposal is thus the inquirer's need to identify proto-knowers. The need to identify people who believe that p comes in only derivatively, via the need to recognize proto-knowers as such. Hence, Craig's genealogy lends succour to the priority of knowledge over belief in this respect. The picture he paints is not one on which humans are first driven to develop the concept of belief and then come to delineate a privileged subset of beliefs as knowledge. Rather, the concern to identify those who proto-know whether p by deploying the concept of proto-knowledge leads the way, trailed by a derivative concern to identify typical but not necessary characteristics of proto-knowers, such as their believing that p .¹¹ In fact, as will become clear once we consider the second part of Craig's genealogy and the process of 'objectivization' it describes, we can even drop the qualification 'proto-', and say simply that the concern to identify knowers leads the way.

6.3 A Genealogy Showing There to Be No Room for Genealogy

This leaves the *no-room-for-genealogy* problem: if the concept of knowledge is primitive and explanatorily basic in the way Williamson suggests (2000, 185), this can seem to exclude any genealogical development towards it. Now on the face of it, there is a straightforward way of reconciling Craigean genealogy and KFE on this front. One can simply point out, as Kusch does (2009b, 90), that the concept of knowledge may be logically primitive without being genealogically primitive: the concept of knowledge may now hold a fundamental place within our conceptual scheme in a way that precludes its factorization into logically prior constituents, but this does not exclude its having developed out of genealogically prior predecessors, and it is compatible with KFE that these predecessors would have been factorizable into constituents. But of course it is open to KFE—and would certainly be in its spirit—to retort that the concept of knowledge is also *genealogically* primitive. And then KFE really would be incompatible with Craigean genealogy.

Or would it? The line I wish to press is that Craig's genealogy *itself* suggests—without becoming entangled in contradiction or forsaking its compatibility with KFE—that the concept of knowledge is genealogically primitive. Looking into the real history of the concept of knowledge, we are unlikely to find an ancestral form of it that is still far from the full-fledged thing, because nothing short of that will do. On this reading, Craig's genealogy lends succour to KFE not only by

¹¹ See Hannon (2019, 43–4) for a different development of the Craigean story according to which belief ends up being a requirement on knowledge.

suggesting that concern with knowledge is prior to concern with belief; it also buttresses and amplifies KFE's claim that the concept of knowledge is logically primitive by explaining *why* it is logically primitive and why we should expect it to be *genealogically* primitive as well.

That the affinities between Craigean genealogy and KFE should run so deep may come as a surprise even to interpreters of Craig who are sympathetic to KFE. Kusch and McKenna (2020, 1068), for example, claim that Craig's genealogy *undercuts* the claim that the concept of knowledge is primitive: by taking the failure of traditional analyses of the concept of knowledge to motivate a genealogical approach to it, they argue, Craig's genealogy undercuts a crucial motivation for KFE, for Williamson takes that same failure to motivate the claim that the concept of knowledge is primitive.

But we need not see these motivations as competing with one another; on the contrary, if the concept of knowledge is taken to be primitive, this just renders all the more pressing the question *why* this primitive notion should have been found useful, since its utility cannot then be explained in terms of the individual utility of its conceptual components. And this is precisely the question that Craig's genealogy answers: it presents the concept of knowledge as *practically basic* and explains why we should expect any human society to operate with something very like the concept of knowledge rather than with the concept of the good informant, because nothing short of the concept of knowledge will do.

There is therefore an important sense in which Craig's genealogy itself shows that there is little room for a genealogy of the concept of knowledge. There remains plenty of room for genealogy *downstream* of the emergence of the concept of knowledge, of course. But if Craig is right, there is hardly any room *upstream* of it. This facet of his account is bound to remain hidden as long as we think of genealogy as mirroring the course of history. On the interpretation of pragmatic genealogy I have given, by contrast, the method is less beholden to the course of history, and there is nothing paradoxical about using a dynamic model to show why we should expect always already to find the concept of knowledge rather than some simpler prototype of it. A pragmatic genealogy may illuminatingly *factor in sequentially* practical pressures that in reality would always already have been at work *simultaneously*. This is a capacity of the method that Craig exploits: the later part of his genealogy factors in needs that in reality do not come *later* at all, but are already at work *alongside* the needs the genealogy starts out from. This implies that earlier stages in Craig's genealogy of the concept of knowledge—the various forms of *proto-knowledge*—are still entirely free of the practical pressures of later stages in a way that no real concept could ever have been. Any real concept would face all of these pressures together, and would thus need to leapfrog Craig's earlier stages to achieve stability in the face of these pressures. This is why the form of the concept that Craig's genealogy issues in must really be the *most primitive serviceable form* of the concept of knowledge.

Nothing in Craig's account excludes our writing histories of how the concept of knowledge was elaborated downstream of that primitive serviceable form. What it suggests is that any such history will likely already *start* with something very like the generic concept of knowledge that Craig ends up with. In reality, the endpoint of Craig's genealogy must be the starting point.

By imaginatively diverging from real history as he does, Craig makes use of a power of the pragmatic genealogical method that we will find on display again in Williams's genealogy of truthfulness in Chapter 7, and it may increase the clarity and plausibility of the interpretation I propose to highlight the shared structure of these two genealogies already here. As we shall see, Williams clearly does not think that, as a matter of actual history, truthfulness first arose in purely instrumental form and only later acquired intrinsic value by coming to be regarded as a virtue. The whole point of his genealogy is rather to show that truthfulness could *only* have been stable insofar as it was valued intrinsically. Similarly, the upshot of Craig's genealogy is not that the concept of knowledge actually first emerged in its prototypical form and then developed into something else. The genealogy represents as arising sequentially what in fact has to arise together, and just as the most primitive form of truthfulness we should actually expect to find already involves intrinsic valuing, so the most primitive form of the concept of knowledge we should actually expect to find is already the 'objectivized' concept Craig reaches in the second part of his genealogy. Craig and Williams both imaginatively diverge from real history in order to convey the same type of counterfactual insight: if X *were* to arise, it *would*, because of the combination of certain practical pressures, likely be driven to develop into Y. This does not carry the implication that X in fact ever existed; on the contrary, it explains why we find Y rather than X, because it shows why, if X ever existed, it would soon have turned into Y.

This is key to understanding the second part of Craig's genealogy, which seeks to explain why we deploy the concept of knowledge rather than of proto-knowledge, and why, if the latter ever existed, it would soon have turned into the former. The concept of knowledge differs from that of proto-knowledge in various ways, most notably in that it picks out more than just good informants: someone might *know* something without being inclined or even able to pass on the information to me. Someone could be useless as a good informant, and thus fail to qualify as a proto-knower, while still qualifying as a knower according to the concept of knowledge we actually have—to take Craig's example, Luigi knows where he buried Mario's body, but he is not telling.¹²

It thus seems that, in contrast to the concept of proto-knowledge, *our* concept of knowledge does not essentially involve the notion of subservience to practical needs, which allows *us* to think what agents at the genealogical stage we are

¹² See Craig (1990, 17, 82).

considering cannot yet think, namely that someone could know something *without* bearing any of the marks of a good informant. This point becomes important when Craigan epistemology is criticized on the grounds that it explicates being a knower in terms of the publicly recognizable properties marking one out as a good informant. Catherine Elgin (2020, 324), for example, points out that the properties regarded in a given society as the marks of a good informant are apt to reflect the pathologies of that society; consequently, to tie the concept of knowledge too closely to these marks leads to epistemic elitism and does an epistemic injustice to those who know their stuff but lack most of the marks of a good informant—Ramanujan was a brilliant mathematician even though his lack of formal training made it difficult for his peers to identify him as such.

This worry has force if we collapse Craig's description of the concept of proto-knowledge into a description of our actual conceptual practice, as Elgin's primary interlocutor, Michael Hannon (2019), in fact tends to do. But if we insist that the genealogical dimension is introduced precisely to prevent us from mistaking a description of the concept of proto-knowledge for a description of our actual conceptual practice, we shall have more room for the thought that while our concept of knowledge is helpfully seen as rooted in the need to identify good informants, so that the concept of *proto*-knowledge initially aligns with the concept of a good informant, there are good reasons why the actual concept of knowledge would have developed beyond that stage. (Indeed, Elgin's critique presupposes that it has, since it turns on the point that someone can be a knower without displaying any of the marks of a good informant.)

Craig's genealogy explains why the concept of proto-knowledge would have turned into something more like our concept of knowledge and thereby emancipated itself from the concept of a good informant. To this end, it factors in pressures Craig takes to be acting on concept-formation more widely, in particular the pressures driving what he calls the process of *objectivization*: the more concept-users resemble social and cooperative creatures with different needs and capacities, the more there are pressures on their concepts to emancipate themselves from their practical origins as private tools answering individual concept-users' needs.¹³ At the beginning of the process of objectivization, concepts are thoroughly indexed to the needs and capacities of specific agents. As objectivization proceeds, this indexation becomes weaker and the concepts less perspectival and more objectivized.

To understand what drives concept-users from purely subjectivist towards more objectivized thought, we need to understand what drives them to distinguish between the *invariable, objective properties of things* that render them useful, and the *variable, subjective needs and capacities* that incite and enable individuals to

¹³ Here I draw on Craig (1990, 82–97) and especially on his Wittgenstein Lectures, available only in German (1993, 81–115).

use those things. It helps to consider, from a first-personal point of view, a primitive form of concept-mongering that lacks such distinctions. At this primitive stage (which Craig ‘slanderously’ imputes to the barnacle, as he puts it), I only have a need and am in want of something that will satisfy it, there and then. What I minimally require is the conceptual capacity to distinguish *holistically* between situations that can satisfy the need and situations that cannot. This is to wield a thoroughly subjectivized concept picking out *whatever can satisfy my present needs, given my current capacities, here and now*.

All but the simplest organisms will be driven beyond this primitive holism, however. I shall be able to exploit many more of the opportunities that the environment affords if I am capable of making more fine-grained distinctions. There are practical pressures on me to distinguish *whatever can satisfy my present needs, given my current capacities, here, now*, from whatever can do so *later*; or *somewhere else*; or *given capacities I anticipate developing* (my energy-reserves are now depleted, but I will soon recover); or *given needs I anticipate having* (I may not be hungry now, but I will soon need food again). I am thus driven to discriminate between these and many more different aspects of opportunity-affording situations and thereby become sensitive to new opportunities afforded by situations I do not currently occupy, but which I might come to occupy.

If I am a *social* and *cooperative* creature, additional advantages will come with sensitivity to opportunity-affording situations that I cannot myself occupy—those available *given someone else’s needs or capacities*, for example, whom I might direct towards such opportunities in the hope of reciprocation. I will likewise be interested in the directions others can give me, and with that interest comes the interest in their operating concepts that are detached from their perspective—just as *they* have an interest in my being able to abstract from my needs and capacities when directing them. Even if I plan to free ride, I need to appreciate their points of view in order to make effective use of them. At the end of this process lie *shared* concepts that are insensitive to the differences between concept-users. They track the objective properties of things that render them suitable to certain uses, irrespective of whether anyone in particular has the need or capacity to use them. Private thinking tools have turned into public ones.

Applied to *proto-knowledge*, these dynamics driving objectivization would similarly lead to the concept’s emancipation from subjective needs. A group of social and cooperative humans that started with this subjectivized concept would be driven to distinguish between the objective grounds of suitability as informants and the subjective needs and capacities to make use of them. Milestones in that process of objectivization include: (i) self-ascription, as individuals come to scrutinize their own qualifications in answer to the question ‘Who knows whether *p*?’; (ii) the direction of third parties to people who might be good informants for them (note that the greater the variety of inquirers an informant has to be suitable for, the more demanding the role of the informant becomes—at the limit, an

informant would be suitable for whomever is asking, whatever their purpose); (iii) the reliance on the identification of good informants by third parties; (iv) group ventures in which individuals need not care about whether they know whether p as long as someone in the group knows.¹⁴

These developments would progressively weaken *proto-knowledge's* indexation to individual perspectives and issue in something close to *knowledge*: the concept of someone who has a sufficient probability of being right about whether p for anyone's purposes, independently of the needs and capacities of any particular inquirer (if the standard for counting as a knower varied wildly from one perspective to the next, knowledge attributions would too rarely allow others to rely on them for a variety of purposes);¹⁵ who may or may not be open, accessible, or intelligible to any particular inquirer; and who may or may not be straightforwardly identifiable to any particular inquirer as satisfying these criteria. Craig concludes that 'the concept of knowing...lies at the objectivized end of the process; we can explain why there is such an end, and why it should be found worth marking in language' (1990, 90–1).

Even if the concept of knowledge did not have much by way of historical predecessors, therefore, Craig's genealogy helps explain why this is so: given the practical exigencies highlighted by the dynamic model of the process of objectivization, nothing short of the concept of knowledge—or at least something very like it—will do. The conviction animating the *no-room-for-genealogy* objection, that the concept of knowledge is basic, thus turns out to be underwritten by Craig's genealogy: the genealogy shows the concept to be practically basic in that any human society would find it hard to get by without it. We are social and cooperative agents who have a need for truths; hence a need for informants; hence a need to identify who proto-knows whether p ; hence a need for the concept of proto-knowledge; hence a need for an objectivized form of the concept of proto-knowledge; hence a need for the concept of knowledge. Only the latter forms an apt response to the combination of needs which Craig derives from uncontroversial facts about generic social and cooperative inquirers. This makes it indeed unlikely that there should have been a gradual historical development towards the concept of knowledge.

6.4 The Art of Our Necessities Is Strange: De-Instrumentalization

What Craig's genealogy tells us goes beyond what it tells us about the concept of knowledge. It highlights far more general dynamics that help explain why concepts

¹⁴ See also Fricker (2010b, 61), Kusch (2011, 9–10), and Hannon (2013, 905–6; 2019, ch. 2).

¹⁵ See Williamson (2005, 101) and Hannon (2013, 916). For discussions of these pressures towards higher standards, see Henderson and Horgan (2015), Grimm (2015), and McGrath (2015).

are driven to shed the traces of their practical origins in individual needs. In doing so, the genealogy enriches and strengthens the methodological framework it operates in—in particular, the *instrumentalism* at the heart of this framework, i.e. the idea that our conceptual practices are illuminatingly approached as *tools* or *techniques* serving practical needs.¹⁶

This kind of instrumentalism about concepts faces two significant hurdles. One is that if we approach concepts as tools originating in the individual concept-user's needs, we should expect most concepts to be tailored to the individual's particular needs and capacities. Yet this is not what we find: most concepts seem to be not personalized in this way, but, as philosophers have been wont to insist, *shared*—as Frege put it, they 'confront everyone in the same way' (1893, xviii).¹⁷

The other hurdle is that the only form of thought that instrumentalism seems capable of explaining is instrumental thought, i.e. thought conforming to the schema 'X is a means by which to satisfy my need for Y'. But we evidently do not view the world solely from such a technical point of view. To reduce human thought to instrumental terms would be to flatten it beyond recognition. Even if our conceptual practices ever were as directly instrumental to the satisfaction of individual needs as instrumentalism represents them as being, they have since been *de-instrumentalized*, largely shedding the traces of their instrumental origins. As King Lear observed: the art of our necessities is strange.

Craig adds to the plausibility of the instrumentalist approach by showing how it can overcome both hurdles. Regarding the first, his genealogy reminds us that while it may well be true that most of our concepts are shared, this is not a conceptual truth—it is a matter of fact calling for explanation; and Craig offers such an explanation in his discussion of objectivization. It explains *why* concepts become shared by becoming progressively less tied to the specificities of particular concept-users. The publicness of concepts, though often treated as a brute fact, thus becomes intelligible as the product of practical demands arising out of two human characteristics: human life is social and cooperative; and concept-users differ in their needs and capacities. Together, these two facts generate pressures on concept-users to operate with concepts that pick out what remains constant across different concept-users' perspectives.

This is not an all-or-nothing matter, as we can conceive of partial overlaps between perspectives as well as of concepts that are only partially objectivized. Some concepts may be more objectivized while others remain tinged with subjectivity. The concepts at work in the natural sciences, especially physics, tend to be more objectivized than those in aesthetics, which remain firmly indexed to

¹⁶ See Brandom (2011) for an overview of this instrumentalist tradition and its roots in the work of Wittgenstein and Rorty.

¹⁷ See Glock (2009, 2010) and Prinz (2002, 14–16, 153). Even Jerry Fodor, for whom concepts are mental particulars, takes it to be a 'non-negotiable condition on a theory of concepts' that concepts be thought of as 'public; they're the sorts of things that lots of people can, and do, share' (1998, 28).

human perception. Craig offers a framework in which to think about why a concept has been objectivized to the extent that it has. Once we think of degrees of objectivization as responses to practical demands, we can reflect on which demands in particular have driven a concept to reach just the degree of objectivization it displays, and what, in a given context, the *point* is of deploying a concept that is objectivized to this degree. In legal contexts, for instance, having concepts that are strongly objectivized serves a point, since much depends on different people using the concept in precisely the same way (think of the efforts that went into objectivizing the concept of a *contract*). In the context of choosing a wine, by contrast, oenological concepts actually serve one better if they are indexed to the drinker's needs and capacities.

As for the second hurdle, Craig's genealogy actually predicts and explains the de-instrumentalization of concepts instead of being derailed by it. His genealogy suggests that there are good instrumental reasons why concepts would come to assume a form that belies their practical origins in the concerns of individual users. This is so particularly if we see Williams (2002) as continuing what Craig started. Both genealogists can be understood as describing different phases on a developmental axis of de-instrumentalization. At the axis' point of origin lies the thoroughly subjectivized and instrumentalized concept indexed to the needs and capacities of the user. At the axis' middle point lies the objectivized but still instrumentalized concept which, while public and no longer indexed to the individual user's needs, is still the concept *of* something instrumental to need-satisfaction and hence indexed to a generic subject's needs. At the axis' endpoint, finally, lies the objectivized and de-instrumentalized concept that is no longer indexed to anyone's needs, because it is no longer the concept of something instrumental to need-satisfaction at all (but, for instance, of something intrinsically valuable).

Williams's genealogical account in effect starts out where Craig leaves off. It shows why, on purely instrumental grounds, dispositions of truthfulness would need to take on more than merely instrumental value for them to possess instrumental value. While Craig's genealogy describes the dynamics driving the development from the point of origin to the middle point, thereby revealing the *instrumental value of objectivized instrumental thought*, Williams's genealogy describes the dynamics driving the development from the middle point to the endpoint, thereby revealing the *instrumental value of non-instrumental thought*. 'It is', as Williams puts it, 'not just at a technical level that we need to understand the world if we are to make a difference to it which we might recognize as an improvement' (2006i, 329–30). We have an instrumental need for non-instrumental thought.

In thinking about the instrumentality of thought, we need to disambiguate between thought that is couched *in* instrumental terms and thought that is instrumental *to* the satisfaction of needs, whatever terms it is couched in. Of

course, thinking in instrumental terms is often instrumental to need-satisfaction. But sometimes, non-instrumental thought will make the better instrument. As we saw in Chapter 3 with the example of loyalty, there are circumstances under which being bloody-minded rather than benefit-minded will reap more benefits—notably when a concept needs to throw new reasons for actions into the balance in order to make a practical difference.

Recognizing the continuity in the dynamics of de-instrumentalization described by Craig and Williams also indicates how Craig's genealogy might be pursued further to describe the development of the concept of knowledge from the middle to the endpoint of the axis of de-instrumentalization, where knowledge is conceived as possessing more than instrumental value.¹⁸ Part of what makes informants *good* informants is that they are disposed to be accurate and sincere; but since, according to Williams (2002, 58–9), instrumental motivations are not sufficient to render these dispositions reliable in the face of the manifold temptations to be inaccurate and insincere, the dispositions to accuracy and sincerity need to come to be regarded as intrinsically valuable dispositions—as *virtues*. If Williams is right, there are instrumental reasons why the notion of goodness encapsulated in the concept of the *good* informant cannot remain simply of the kind involved in being a good knife—the goodness of a tool suited to its purpose. Rather, the concept of the good informant, and thus of proto-knowledge, needs to involve intrinsic valuation in order to sustain the institution of information pooling that gives the concept its point. This insight might be used to explain why knowledge has been deemed superior in value to true belief, and why achieving knowledge has been thought to involve the exercise of virtue.¹⁹

The second hurdle for instrumentalism about concepts can thus be overcome using the idea of the instrumentality of non-instrumental thought. By giving instrumentalist approaches to concepts a genealogical dimension, Craig and Williams enable us to grasp the practical origins of concepts in the needs of concept-users while also explaining why these concepts have come so far from their instrumental origins.

6.5 Assessing and Synthesizing Competing Accounts of Concepts

Finally, an important contribution of Craig's account is that he harnesses pragmatic genealogy's ability to help us assess and reconcile competing accounts of concepts. He reminds us that our concepts are, or should be, keyed to the world we

¹⁸ See Kusch (2009b, 74–6).

¹⁹ See Kvanvig (2003), Sosa (2007), Kusch (2009b), Fricker (2009), Pritchard (2012), and Hannon (2019, 63).

live in, so that understanding what concepts would serve us well in that world can give us a sense of what our concepts are likely to be, which in turn allows us to assess the plausibility and compatibility of competing accounts of them.

Craig's own genealogy informs our take on extant conceptual analyses of knowledge in several ways. First, it renders us suspicious of any account of the concept which makes it so demanding that its extension would effectively be empty. A concept that put the threshold for someone to count as knowing that *p* too high would be utterly pointless, since it would simply fail to get a grip on the world we inhabit—nothing would ever satisfy the concept. From a practical perspective, it is clear that the concept should at least *originally* permit some form of knowledge, and that the concept would have lost that original point if it now no longer did so.²⁰ Miranda Fricker (2008, 48) maintains that we can enlist this insight into the genealogical primacy of knowledge-permitting situations against scepticism about the possibility of knowledge. A concept of knowledge that was as demanding as the concepts involved in formulating sceptical challenges would leave inquirers seeking good informants empty-handed. The state of nature, she writes, 'explains the commonsense idea that no one can basically be a sceptic. They must be inquirer first, and sceptic second; someone committed to the practical possibility of knowledge first, and committed to undermining that possibility second' (2008, 46). As she goes on to acknowledge, however, Craig's genealogy only shows that the concept of knowledge must *originally* have been knowledge-permitting. This does not by itself exclude its subsequent development into a concept allowing us to mount sceptical challenges.²¹

Second, Craig's genealogy accounts for the variety of conditions that have been advanced to specify what must be the case for the concept of knowledge to apply. All the familiar candidates—standing in the right causal relation to the state of affairs in question, having proven reliable in the past, or being able to offer justifications—are shown to have a common practical origin in the inquirer's need to *recognize* the good informant as such.

Third, the genealogy accounts for the vexation surrounding attempts to privilege any one of these conditions at the expense of the others, because it shows that there is no practical reason for users of the concept to do so. Each of these conditions does its bit by rendering inquirers sensitive to circumstances in which informants typically are good informants, and the wider the net is cast, the greater the catch.

Fourth, the genealogy suggests that the conditions are not necessary, but typical conditions: conditions worth tracking because doing so *typically* pays off. From a

²⁰ See Hannon (2019, ch. 3) for a discussion of how Craig's genealogy informs our view of the threshold question.

²¹ Fricker then goes one step further, however, claiming that 'even the sceptic cannot escape the cognitive functionality of the origin, for that scenario is still with us, at the core of what it is for us to know' (2008, 49). I discuss whether this further inference is licensed in Chapters 8 and 9.

practical point of view, what we want are concepts that are maximally cost-effective: they should strike an optimal balance between the costs and benefits of using them in the kind of world in which they are deployed (this is what it is for them to be *keyed* to the world). One notable way in which a concept can be cheap to use is by tracking properties that are readily accessible; if it tracks properties that it takes a lot of time and energy to get at, using the concept will come at a higher cost. On the benefit side of the balance sheet, a concept earns its keep by serving its point *in practice*, which usually means over a series of confrontations with real-life cases; the idea of a concept whose application conditions, when satisfied, virtually *guarantee* that the concept will serve its point sounds enticing, and evidently holds a certain fascination for philosophers; but aiming at perfection is the enemy of the good. A concept whose application conditions are such that their satisfaction virtually guarantees the presence of what the concept is supposed to help us track, but which makes it very costly, in terms of the time and effort required, to determine whether the conditions are in fact satisfied, may well prove *less* valuable in the long run than a concept that is riskier, in that the satisfaction of its application conditions only makes it *probable* that its point will be served, but that is cheaper and easier to deploy. By these measures, Craig's model suggests that tracking the wide range of properties that *typically* mark a good informant is a better bet than tracking the narrow set of properties that a good informant *necessarily* has.

Fifth, the genealogy calls into question the ambition to spell out criteria that consistently pick out knowledge even under the most freakish imaginary circumstances. By definition, we encounter freakish cases only rarely, if at all, so that they will have left less of a mark on the concept. As Craig himself notes, moreover, the unpredictability and variability of freakish circumstances means that to 'try to make a practice of detecting freakish cases' with near-infallibility would likely involve 'high costs in time and energy' (1990, 15) which are unlikely to pay off, not least because the cases are rare. Hence, he concludes that 'it must be the standard or prototypical case at which the inquirer's strategy is directed' (1990, 15). Of course, freakish circumstances may also play into inquirers' hands: 'a purpose may be achievable in unusual ways—factors which would usually frustrate it may, if other features of the situation are exceptional, do no damage, factors which are usually vital may, abnormally, be dispensible' (1990, 15). But this only reinforces the conclusion that an inquirer had better be guided by a range of typical conditions rather than a narrow set of necessary conditions, because this provides a flexibility that allows the concept plially to accommodate varying circumstances. What philosophers should be aiming to spell out in giving an account of the concept of knowledge, therefore, are not conditions obtaining even across freakish cases, but conditions enshrining 'the features that effect realization of the purpose when things are going on as they nearly always do' (1990, 15).

More generally, it follows from the idea that we should expect and want our concepts to be keyed to the world we live in that considering their applicability to wildly counterfactual worlds is unlikely to tell us much about our concepts and their merits. Take, for example, concepts of epistemic evaluation such as those we employ in calling someone *irrational* or *unreasonable*. Let us hypothesize that these concepts perform the function of reforming bad truth-trackers by taking as proxies such indicator properties as the inability to offer a justification for one's belief. If the world happened to be full of unwittingly reliable clairvoyants, tracking the inability to offer justifications would no longer typically pick out bad truth-trackers. In such a world, our concepts of epistemic evaluation would, *ex hypothesi*, cease to function well, because they would be premised on heuristic assumptions that did not generally hold true. But given that our world does not contain unwittingly reliable clairvoyants, it is no wonder that our concepts are such that they lead us to find epistemic fault with these hypothetical characters. Our concepts are keyed to a world in which truth-tracking and reason-giving tend to go together. That they go together is a contingent fact about our world. But the fact that our concepts would not serve us well *if* that contingent fact did not obtain in no way detracts from the functionality of our concepts in *our* world.

Likewise, we should expect the functionality of concepts in our world to admit of counterexamples: cases where the conditions tracked by the concept are given but applying the concept fails to serve anyone's needs or interests. As we saw in Chapter 2, explanations in terms of typical conditions are more tolerant of exceptions than explanations in terms of necessary conditions. This fact does not reflect a lack of ambition or a willingness to settle for less than complete accounts of concepts. There is a compelling metaphilosophical rationale for it. A concept need not track necessary conditions in order to be serviceable. In fact, it would manifest a strange impoverishment in our conceptual resources if they only rendered us sensitive to perfect regularities—it would mean forsaking all opportunity to latch on to less than perfect regularities. We live in a world in which things tend to come in continua of variations on prototypes rather than in distinct series of perfect clones. In such a world, concepts tracking necessary conditions—such as the concept of a triangle in geometry—have their uses, but we should expect the vast majority of our concepts to be more attuned to protean variety than to rigid uniformity. Where there is no accessible set of properties waiting to be tracked that will *guarantee* the fulfilment of a concept's function, the obvious solution is to *build the variety right into the concept*: to track the variety of conditions under which the concept typically, though not necessarily, fulfils its function. This increases the concept's functionality by rendering it more tolerant of variation, so that it more than makes up in flexibility what it lacks in precision—just as we might prefer a mechanism that can accommodate changes in its environment as it performs a task over a mechanism whose capacities break down as soon as variation is introduced, even if the former is less clinical in its

execution than the latter. *Variation tolerance*, as one might call it, is an important desideratum on concepts in the kind of world we inhabit.

The method of pragmatic genealogy is well suited to bringing out the formative influence of such practical considerations on our concepts, and by showing how these considerations account for the internal diversity within our concepts, the method also indicates ways to *synthesize* competing accounts of concepts. Craig's project is the best example. The 'conceptual synthesis' of his title is doubly appropriate: not only does he offer an alternative to conceptual analysis by looking at how something like our concept of knowledge might be *put together* instead of analysing the concept we actually have into its constituents; he also synthesizes various accounts of the concept. Treating extant accounts of the concept as data, he shows that we do not need to decide between them. Each analysis latches onto some of the criteria that really do guide the concept's application. The competing accounts only make the mistake of treating these criteria as exhaustive or as necessary conditions. Craig offers a synthesis by presenting the various accounts as partial but complementary descriptions of an internally diverse conceptual practice. The debate between internalists and externalists, for example, which turns on whether people only know that *p* if they are able to articulate the justification for *p*, loses much of its point in light of Craig's account (1990, ch. 8). As Fricker puts it: 'The picture in the State of Nature is fundamentally externalist—what matters is simply that good informants come out with the truth—but then we quickly come to see the origin of internalist intuitions about knowledge' (2008, 42). We come to see that informants' ability to give justifications is typically a good indicator of their being good informants, which explains why our concept of knowledge would track this ability and hence also why the internalist analysis latches onto a real feature of our concept. Kusch and McKenna (2020, 1060–1) point out further competing theoretical intuitions concerning the concept of knowledge that Craig vindicates and reconciles: that one can know that *p* without believing that *p* (1990, 15–16); that counterfactuals play a role (1990, ch. 3); that causal relations play a role (1990, ch. 4); that the reliability of one's methods plays a role (1990, ch. 4); that standards vary with context (1990, ch. 12); and, last but not least, that all analyses of knowledge in terms of necessary and sufficient conditions admit of counterexamples (1990, ch. 6).

The promise held out by pragmatic genealogy is thus that stand-offs between competing accounts of a conceptual practice might be resolved by showing that each correctly identifies *some* property worth tracking given the concept's function, but that in light of that function, we should expect the concept to track *all* these properties, because they form the collection of conditions under which the concept typically serves its point.

In this chapter, I have examined three reasons to think that Craig's genealogy of the concept of knowledge is incompatible with KFE, and found that far from being

incompatible with it, Craig's genealogy actually supports it. Even if the concept of knowledge did not in fact have much by the way of historical predecessors, we can still give an imaginary genealogy to help explain why it predominantly exists in the form we know. I have then argued that Craig's contribution goes beyond what it tells us about the concept of knowledge, bearing on such questions as why concepts are shared, why particular concepts have been objectivized the way they have, and why concepts are internally diverse. Along the way, I have claimed that there is an underexplored continuity between Craig's and Williams's genealogies: they both describe dynamics of de-instrumentalization that combine to yield a powerful framework within which to think about concepts and their relations to our needs. But much more remains to be said about Williams's continuation of Craig's project. It is therefore to Williams that we now turn.