MARTIN KUSCH

KNOWLEDGE AND CERTAINTIES IN THE EPISTEMIC STATE OF NATURE

"… it is important to imagine a
language in which *our* concept
'knowledge' does not exist."
(Wittgenstein)

ABSTRACT

This paper seeks to defend, develop, and revise Edward Craig's "genealogy of knowledge". The paper first develops the suggestion that Craig's project is naturally thought of as an important instance of "social cognitive ecology". It then introduces the genealogy of knowledge and some of its main problems and weaknesses, suggesting that these are best taken as challenges for further work rather than as refutations. The central sections of the paper conduct a critical dialogue between Craig's theory and Wittgenstein's claim – familiar from *On Certainty* – that common-sense certainties cannot be known. It turns out that Craig's distinction between different stages in the development of our concept of knowledge can illuminate and make plausible Wittgenstein's claim. But it can do so only if Craig's traditional commitment to a central "core" in our concept of knowledge is replaced with the idea of *knowledge* as a family-resemblance concept.

INTRODUCTION

In this paper I want to push further a line of inquiry that I began in "Testimony and the Value of Knowledge" (2009): to explain important features of our epistemic concepts using, as my framework, a substantially modified version of Edward Craig's functional-historical method of analysis, known as "conceptual synthesis", "state-of-nature epistemology", or "genealogy of knowledge" (Craig 1990; cf. Williams 2002). In the earlier paper as well as here, I argue for the need to modify Craig's approach by bringing his work into a critical dialogue with authors who have pursued related projects. In what follows, Craig's main interlocutor will be Ludwig Wittgenstein, especially the author of the late notebooks published under the title *On Certainty* (1969). On the one hand, I shall use Craig's ideas

Episteme 8.1 (2011): 6–23
DOI: 10.3366/epi.2011.0003

www.eupjournals.com/epi

to make sense of Wittgenstein's notorious claim that Moorean common-sense certainties are not "knowledge-apt" (to coin an expression). On the other hand, I shall show that Craig's theory can be considerably strengthened by combining it with a Wittgensteinian conception of "family-resemblance concepts".

"THE CONCEPT OF KNOWLEDGE IN OUR SOCIAL COGNITIVE ECOLOGY"

Before I get to my main topic, however, I want to reflect briefly on the title of the 2010 Edinburgh *Episteme* conference to which this paper was a contribution. I wish to explain how I understand the title and why I find it a suitable heading for my preoccupations with Craig and Wittgenstein.

I shall make my way from the back to the front. *Ecology* is the study of organisms and their surrounding environment, both living and non-living. *Cognitive ecology* is the investigation into how organisms obtain information about their surroundings. So far, so familiar. But "social cognitive ecology" does not yet have an established use. We need to decide on its content. I submit that the "social" in "social cognitive ecology" qualifies both the organisms and the environment. Adding a bit of flowery language, this gives us the characterization of the central aspects of the new field as: the inquiry into how highly gregarious, deeply interdependent cognizers (like us) obtain information about their natural and social environment.

Unfortunately, this characterization does not yet sufficiently distinguish social cognitive ecology from other fields like social epistemology or the sociology of knowledge. To achieve a greater product distinction, we need to return once more to ecology itself. We need to remember that ecology is a natural science, not an investigation of our normative intuitions. Accordingly, social cognitive ecology should be *descriptive-explanatory* rather than *normative*, and it should, wherever possible, reckon with the results of the natural and social sciences. Moreover, ecology and evolutionary biology are conventionally considered to be "sister disciplines"; indeed, the most important journal in ecology is called "*Trends in Ecology and Evolution*". This link makes it imperative to tie ecology to concerns with change and development, that is, with real and imagined history.

Social cognitive ecology can investigate a wide range of phenomena; one obvious set of phenomena – as the conference title is right to signal – are our epistemic concepts. A social cognitive ecology of our epistemic concepts should explain why, in and through their historical change and development, our concepts of belief, knowledge, justification, or wisdom (amongst others) have received their current shape.

Who can lay claim to being a "social-cognitive ecologist" in my sense? As I see it, there are four candidates:

(1) the genealogy of knowledge,
(2) the later Wittgenstein,

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

(3) the sociology of knowledge (e.g. Barnes, Bloor, & Henry 1996), and
(4) "historical epistemology" (e.g. Daston and Galison 2007).

(Historical epistemology is a subfield of the history of science interested in the emergence and change of categories like "objectivity", "proof", or "observation" in the sciences.)

I believe that these four projects supplement and (potentially) correct and challenge one another. Since this is my view, I have elsewhere critically related (1) and (3), and then again (1) and (4) (Kusch 2009, 2011). As already announced, in this paper the central focus will be on a critical dialogue between (1) and (2).

CRAIG'S GENEALOGY – THE BASIC IDEA

I shall now introduce the main ideas of Craig's 1990 book *Knowledge and the State of Nature*. To facilitate my exposition, from here on, I shall write *concepts* in italics, 'words' in inverted commas, and leave elements of the extension unmarked. Every occurrence of "to know" or "knowledge" is to be read as "to know (and its cognates)".

The central proposal of *Knowledge and the State of Nature* is to replace traditional conceptual *analysis* with conceptual *synthesis*; that is, with a hypothetical historical narrative of the process in which *knowledge* was put together. This narrative is constrained by two ideas.

The first idea is that of an epistemic "state of nature", that is, of an imaginary early social community composed of language-using human beings who are co-operative though not kin, and whose conceptual and reflective powers are somewhat weaker than our own. The genealogical narrative must make intelligible why these creatures found it useful or valuable to introduce an ancestor of our concept of *knowledge*. I shall call this ancestor concept "protoknowledge". The second constraint is developmental: the genealogical narrative must explain why the ancestor concept was eventually replaced by our concept. I shall explain the two stages in turn.

In the state of nature, individuals depend upon one another for information. Distinguish between the roles of "inquirer" and "informant". The inquirer needs information that she is currently unable to directly obtain herself; the informant offers such information. Inquirers must be able to separate good from bad informants. And it is natural to assume that meeting this need will involve concepts. Assume that *protoknower* is the central conceptual tool for dealing with this problem. Which conceptual components should *protoknower* contain? What should we hypothesize our imaginary ancestors to want this concept for? Craig's answers are that our ancestors want this concept as a tag for good informants and that the concept *protoknower* (*whether p*) comprises the following elements: (i) being as likely to be right about *p* as the inquirer's current needs require; (ii) being honest; (iii) being able to make the inquirer believe that *p*; (iv) being accessible to the

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

informant here and now; (v) being understandable to the inquirer; and (vi) being detectable as a good informant concerning *p* by the inquirer. To elaborate briefly only on the last point, the inquirer needs to find properties (=*X*), 'indicator-properties', that she is able to detect and that correlate closely and in a law-like fashion with holding a true belief, or telling the truth, as to whether *p* (1990, 25, 135). 'Being at the top of a tree' might be such a property *X* for some inquirers in the state of nature when *p* is the proposition that a tiger is approaching the village. Usually more than one *X* will be involved. The properties that make Fred a medical protoknower are not one but many. Craig is adamant that (i) to (vi) are not necessary and sufficient conditions. While all of these elements are present in prototypical situations, the concept has a use even when some of the elements are missing. Finally, *protoknowledge* differs from *knowledge* in that: (a) only the former is closely tied to testimony; (b) *protoknowledge* is not a fully public concept insofar as it is indexed to the capacities and needs of specific inquirers (90); (c) *protoknowledge* can be ascribed only to others but not to oneself; and (d) protoknowledge is not undermined by accident or luck: users of *protoknowledge* lack the intellectual sophistication to distinguish between accidental and non-accidental fulfillment of the conditions of *protoknowledge*.

This brings us to the second half of the genealogical just-so story: the hypothetical social-historical narrative that takes us from *protoknowledge* to *knowledge*. Craig speaks of this development as a process of "objectivisation" of *protoknowledge*. Key steps in this objectivisation are the following. First, *protoknowledge* comes to be used in self-ascription. In response to the question "who knows whether *p*?" group members start to investigate their own indicator-properties. Second, inquirers begin to recommend informants to others. This can be done in a helpful manner only if the perspectival or indexical character of *protoknowledge* is weakened. The recommended informant must be good in the eyes of both the recommender and the recipient of the recommendation. Further movement in this dimension – recommending an informant to ever more inquirers – makes protoknowledge increasingly harder to get. The endpoint is the idea of "someone who is a good informant as to whether *p* whatever the particular circumstances of the inquirer.... That means someone with a very high degree of reliability, someone who is very likely to be right – for he must be acceptable even to a very demanding inquirer" (91). And a very demanding inquirer will not accept epistemic luck or accident. Third, inquirers begin to use "being recommended" as an indicator-property. This move dilutes the original detectability requirement. Inquirers begin calling someone a "protoknower" even when none of the original "natural" indicator-properties is in sight. Fourth, in the context of group action, inquirers cease to care whether the needed information is accessible to them as individuals; they are satisfied if it is accessible to someone in the group. As a result, they will speak of protoknowledge even outside the context of testimony. The process of objectivisation ends up with our concept of *knowledge*: "The concept of knowing... lies at the objectivised end of the process; we can

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

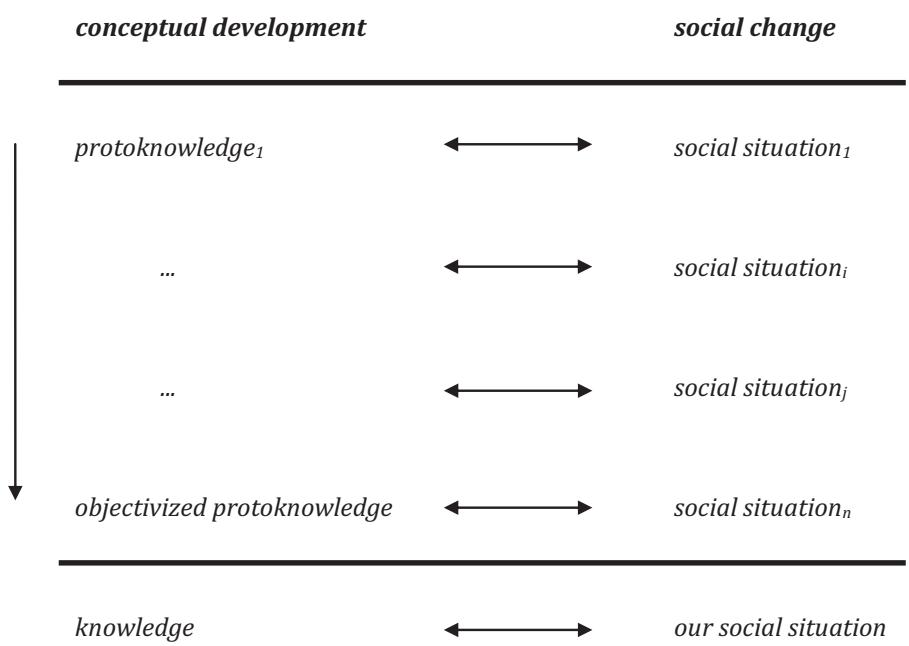

Figure 1. The genealogy of knowledge as a diagram.

explain why there is such an end, and why it should be found worth marking in language" (90–1).

In my view the most important aspect of objectivisation – an aspect not much highlighted by Craig himself – is the thought that the development in the conceptual realm is tied to, and explained by, changes in patterns of interaction in the group. Perhaps this point can be captured by saying that concepts (or stages) of protoknowledge are "salient solutions" (cf. Schelling 1960) to different social coordination problems. The graph above (figure 1) highlights this idea and sums up my exposition of Craig's work.

CRAIG'S GENEALOGY – DESIDERATA AND GAPS

In order to evaluate Craig's theory, it helps to distinguish between four criteria that a successful genealogy should be able to meet. First, it should deliver *objectivised protoknowledge* that is extensionally and intensionally equivalent (or at least close) to our *knowledge*. Second, the conceptual synthesis (the path from *protoknowledge* to *objectivised protoknowledge*) should be plausible both philosophically and when judged by the results of, say, historical linguistics. Third, the claims about the right-hand side of figure 1, that is, the claims about social patterns and their development, should be convincing in light of social theory, cognitive psychology, and related fields. And finally, fourth, the conceptual development on the left of figure 1 should match up plausibly with the social patterns on the right.

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

Critics have registered complaints in all four areas. I shall here briefly note the most important ones, without attempting an assessment of their merit at this point. I begin with the left-hand side of figure 1, combining the first two criteria. Timothy Williamson objects to Craig's very starting point. As Williamson sees it, Craig goes wrong when he stipulates that inquirers want "true beliefs about our environment, as though this were somehow more basic than our need for knowledge of the environment" (2000, 31). In other words, Williamson sees no logical space for a conceptual development towards *knowledge*.

Christoph Kelp (this issue) raises two important objections. He begins by pointing out that the purpose or interest relativity of *(objectivised) protoknowledge* is much greater than even authors sympathetic to purpose relativity (e.g. Hawthorne 2003 and Stanley 2005) allow for in the case of *knowledge*. Moreover, Craigian objectivisation results in the conception of a knower "with a high degree of reliability, someone who is very likely to be right – for he must be acceptable even to a very demanding inquirer" (Craig 1990, 91). Kelp notes that this characterisation befits "high-standard classical invariantism". Unfortunately, Craig offers little in support of this view. Why not hold instead, say, that objectivisation leads merely to the demand that the knower be reliable so as to satisfy *ordinary* concerns?

Turning to the right-hand side of figure 1, here the most important commentators have been Bernard Williams and Miranda Fricker. Williams (2002) questions Craig's assumption that the inhabitants of the epistemic state of nature are cooperative. As Williams sees it, cooperation cannot be taken for granted in this way. Hence Williams insists on, and develops, an account of how information-sharing is possible as a social institution and collective good. Williams is also unhappy about the absence of "real" history in Craig's story. Put differently, Williams believes that Craig's "imaginary genealogy" needs to be complemented by "real genealogy", that is, by an engagement with historical and cultural contingent realities. Williams insists that in moving from imaginary to real genealogy, we do not leave philosophy behind: "… philosophy cannot be too pure if it really wants to do what it sets out to do" (39). Fricker finds both Craig's and Williams's characterizations of the epistemic state of nature insufficiently *political*: neither of them says anything about the effects of social categorization upon one's status as a good or bad informant (Fricker 2007).

In some of my own previous work I have commented critically on the fourth demand, that is, the need to tie together conceptual and social developments. Going beyond Craig, I have proposed that *protoknowledge* attributions might have a role in maintaining the collective good of information sharing: they honor and thus reward the good informant. That is to say, such attributions acknowledge that the informant has contributed to the collective good by having made "investigative investments". I also put forward the thought that we should make Craig's scenario more realistic by considering not only the conceptual needs of the inquirer (in distinguishing between good and bad informants) but also the conceptual needs of the informants (in separating out good and bad inquirers).

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

Assuming the criticisms are roughly on target, where would this leave Craig's genealogical proposal? Should we *reject*, *replace*, or *reform* it? It is hard to see how Williamson could opt for anything less than rejection; if knowledge is primitive in the way he proposes, there is little space for any genealogy. Kelp inclines more towards replacement. He wishes to change the focus of the genealogy of knowledge from good informants to the point "when agents may adequately terminate inquiry into a given question." Williams, Fricker, and I are reformers: while we take the worries and criticisms seriously, we try to treat them as *challenges* for, rather than refutations of, the genealogical program. That is to say, we seek to improve genealogy in ways that enable it to circumvent the objections. And we aim to do so in ways that are illuminating and explanatory, and that, whenever possible, draw on other (scientific) fields. In this paper I shall try to do so especially regarding the concerns about the right-hand side of figure 1 and with respect to the links between the two sides. Responding to Kelp's and Williamson's objections will have to be left to future work.

WITTGENSTEIN AND THE GENEALOGY OF KNOWLEDGE – GENERAL ISSUES

In that reformist spirit I now turn to the dialogue between genealogy and the later Wittgenstein. I begin by noting some general similarities and differences between the two authors.

First, Craig and Wittgenstein agree on the need for "just-so" stories about (natural) pre-history. Wittgenstein writes, for example, "… we are not engaged in natural science, and not even in natural history – since we can also surely provide fictitious natural history for our purposes" (1953, II, xii). Intriguingly, he also points out that "it is important to imagine a language in which *our* concept 'knowledge' does not exist" (1969, §562). Genealogy does just that.

Second, Wittgenstein too treats conceptual and social change as inseparable. After all, he urges us to focus on how language is used, that is, on the function of different words and sentences in specific types of social interactions. Changes in patterns of social interaction are internally related to changes in our concepts (Winch 1958; Bloor 1997).

And third, Craig and Wittgenstein both hand-wave in the direction of *knowledge* as a family-resemblance concept. As already mentioned, Craig rejects the notion that *knowledge* can be captured in terms of necessary and sufficient conditions. To this we can now add that Craig does not seek to assimilate all uses of "knowledge" to his "good informant" account. In discussing expressions like "knowing how", "knowing people", or "knowing places", he shows himself open to the possibility of an non-reducible variety of uses (1990, 140–61). I shall give some evidence of Wittgenstein's views later in this paper.

There are also some important differences between genealogy and the later Wittgenstein. One is the extent to which Wittgenstein stresses the contingency

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

of our possessing certain concepts. In Craig this consideration plays little role. Wittgenstein writes:

> It is often useful, in order to help clarify a philosophical problem, to imagine the historical development… as quite different from what it actually was. If it had been different, no one would have had the idea of saying what is actually said. (1978, III, 81)

Moreover, Craig does not present his genealogy as a principled criticism of mainstream epistemology. Wittgenstein is much more critical of what he sees as the mainstream epistemologists' tendency to "sublimate" our ordinary epistemic concepts:

> Some will insist that what I say about the concept of knowledge [e.g. that 'knowledge' is not always factive] is just irrelevant. They will admit that the philosophers' understanding of the concept of knowledge is different from the everyday understanding. But they will go on to argue that the philosophers' concept is more important than the ordinary and not very interesting concept…. However, the philosophical concept was extracted from the ordinary concept via all sorts of misunderstandings… (2000, 232/679/288; 11.1.1948)

CERTAINTIES IN THE EPISTEMIC STATE OF NATURE

The significance of the last-mentioned difference in particular will emerge in due course as I develop my main thesis that Craigian genealogy can be used to throw new light on Wittgenstein's claim that common-sense certainties are not knowledge-apt.

The common-sense certainties of *On Certainty* are propositions like

> The earth has existed since long before my birth.
> Every human being has two parents.
> Water boils at 100° C.
> The earth is round.
> I am a human being.

To be precise, these "Moorean" propositions are common-sense certainties only when they are taken for granted; the moment their truth is in question, they cease to be certainties.

Wittgenstein says many surprising things about common-sense certainties, for instance, that they are neither true nor false, neither justified nor unjustified, or neither empirical nor normative. In this paper I want to concentrate on making genealogical sense of this claim: "I should like to say: Moore does not know what he asserts he knows…" (1969, §151). This statement must be understood in the context of Wittgenstein's argument that Moore – in his "A Defense of Common Sense" and "Proof of an External World" – is *sublimating* common-sense certainties: Moore decontextualises common-sense certainties and treats them as the most certain knowledge, thereby losing sight of the social coordination

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

problems for which common-sense certainties are (part of) the salient solution; Moore misconstrues common-sense certainties as universally valid intuitions, ignoring the ways they might vary with different cultures; and Moore thinks of common sense as a "mind-set" rather than as a set of socially shared assumptions.

To see how Craig's genealogy can be used to show that common-sense certainties are not knowledge-apt, we need to return once more to the elements that get "synthesised" into the concept of *protoknower*. Protoknowers are informants who fit the needs of the inquirer; that are honest towards the inquirer; that are able to convince him or her; that are understandable to the inquirer; and that are detectable by the inquirer (on the basis of "indicator properties"). For our purposes it suffices to focus on the first condition. It obviously is shorthand for a number of elements:

(1) The information is relevant to the inquirer's current and/or likely future situations.
(2) The information is new to the inquirer.
(3) The inquirer cannot obtain the information with minimal effort herself.
(4) The inquirer's ability to reliably obtain the information does not much exceed that of the informant.
(5) Obtaining the information is (for the informant) an achievement.

An utterance of a common-sense certainty does not fulfill any of these five conditions: common-sense certainties are too basic to be relevant to inquirers' plans and deliberations; since everyone knows them "always already", they are no news; the issue of effort does not arise in their case; there are no significant differences in reliability regarding them; and "reporting" them is not an achievement. Conclusion: common-sense certainties are not "protoknowledge-apt".

That common-sense certainties are not protoknowledge-apt does not mean that they may not have important functions in an epistemic state of nature or in our institution of information-sharing. Perhaps it is most natural – in light of what Wittgenstein himself says in *On Certainty* – to think of them as "ultimate negative indicator properties". Someone who explicitly or implicitly denies common-sense certainties thereby proves himself unable to play any useful role as either informant or inquirer. Such people do not meet the condition of being "rational", or "*vernünftig*":

> So it might be said: "The rational man believes: that the earth has been there since long before his birth, that his life has been spent on the surface of the earth, or near it, that he has never, for example, been to the moon, that he has a nervous system and various innards like other people, etc., etc." (1969, §327)

Understood in this way, common-sense certainties might also play an important role in the social-political processes of exclusion that Fricker's work has made so salient. After all, many CSCs concern bio-social identities:

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

I am a man … (§79)
I know that my name is L.W. (§328)

The proposal according to which common-sense certainties are ultimate negative indicator properties and, as such, criteria of rationality also throws some light on Wittgenstein's claim that common-sense certainties cannot be successfully or meaningfully doubted. After all, to doubt the *criteria of reason* is to place oneself on the side of unreason – at least this is so in a state of nature in which philosophical reflection is not yet a major pastime. Moreover, to doubt reason is to doubt the *criteria of membership* in our community – and this is to place oneself outside of it. Hence doubting common-sense certainties is both epistemically and socially explosive: sceptics refuse to accept the "authorities" that one "must recognise … in order to make judgements at all" (§493); they opt out of our "community which is bound together by science and education" (§298).

CAN WE KNOW CERTAINTIES?

To show that common-sense certainties cannot be *proto-known* is not of course to make plausible that they cannot be *known*. So far I have only shown the former, not the latter. I suspect that most of my readers will be doubtful as to whether the latter can be achieved at all. After all, the intuition that common-sense certainties can be known is strong (at least amongst epistemologists), and there seem to be good pragmatic, Gricean reasons for explaining Wittgenstein's sense that claiming knowledge of common-sense certainties is pointless: to "report" common-sense certainties violates the "maxim of quantity" (make your contribution as informative as it required) and the "maxim of relation" (be relevant) (Grice 1989).

My response will consist of two *partial vindications*. First, I shall vindicate the majority of epistemologists who insist that common-sense certainties are knowledge-apt. I shall support the underlying intuition by proposing a conceptual-cum-social development in and through which *knowledge* becomes detached from the ideas of credit, honouring, social status, achievement, and relevance. Obviously, it is only once *knowledge* loses these features that it becomes possible to think of common-sense certainties as knowledge-apt. Having partially vindicated the epistemologists, I then hurry to Wittgenstein's partial defence: I argue that he has a point – and a point that owes nothing to Grice – when he denies that common-sense certainties are knowledge-apt. The result of both vindications will then prove to depend on the idea that *knowledge* is a family-resemblance concept.

HOW DID COMMON-SENSE CERTAINTIES COME TO BE KNOWLEDGE-APT?

The question is easy to answer: *(proto-)knowledge* must lose the elements of (1) (immediate) situational relevance, of (2) an achievement worthy of honor or credit,

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

and of (3) a social status. A genealogical story of these processes might take the following form.

(1) As Craig rightly emphasizes, objectivisation means that *(proto-)knowledge* becomes increasingly hard to get; it must be relevant to more and more inquirers in more and different situations. This requires ever costlier investigative investments. These costs are likely to push people in the state of nature to develop systems of information storage. Once the state-of-nature equivalent of a databank is in place, information need not be obtained from a first-hand source every time anew; the information can be kept for future reference. The important point is this: information stored in this way is no longer immediately relevant; until it is needed, its relevance is merely potential or virtual. Indeed, one might say that until the information is retrieved in a specific situation, it is not relevant at all. It is likely that this development leaves its marks on *(proto-)knowledge*: information need not be relevant to qualify.

(2) Recall that *protoknowledge* attributions are honoring so as to motivate informants to make investigative investments (e.g. climb up on the tree to tell us whether the tiger is approaching the village). But the need for this kind of honoring diminishes as two things happen in complex societies: testimonial chains grow longer, and different people can be asked for the same information. In long testimonial chains only very few subjects make substantive investigative investments; most act as mere conduits. Given the very small effort involved in acting as a conduit, people do not need much encouraging. Moreover, if I can get the needed information from many different informants, there is no need for me to spend too much effort honoring any particular one.

(3) Although (1) and (2) suffice to explain why common-sense certainties might have become knowledge-apt, it seems appropriate to add a third – if only because it is important to Wittgenstein. The third shift is from *(proto-)knower* as a *social status* to *knowing as the mental state* needed (at least as a necessary condition) for possessing good information. This change may well not be linked to a particular social development; social scientists tell us that shifts from "status to state" are a common phenomenon in humans of all kinds of societies (Barnes 2000). (Think of how tempted we are to think of coins and bills as intrinsically valuable.) The important point for the present context is that such development has the effect of further detaching *knowledge* attributions from their original social coordination context.

Note that I am not treating this "reification" as a pathology but as a feature of human psychology that leaves its marks on our concepts. Wittgenstein takes a different view; he thinks that conceiving of knowing as mental state is part and parcel of philosophers' misunderstanding of ordinary uses of "knowledge":

> All psychological terms merely distract us from the thing that really matters. (1969, §459)

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

> Knowledge in mathematics: Here one has to keep reminding oneself of the unimportance of the 'inner process' or 'state' … (§38)
> We are asking ourselves: what do we do with a statement 'I *know* …'? For it's not a question of mental processes or mental states. (§231)

Be this as it may, the important result of the above consideration is the following: once *(proto-)knowledge* has lost the features of relevance and achievement, and once the process of reification has occurred, it is perfectly correct for me to ascribe to myself the *knowledge* that the earth has existed since long before my birth.

THE RECEIVED VIEW

In this section I begin my multiple-step attempt to partially defend Wittgenstein's insistence that common-sense certainties are not knowledge-apt. The first step is the sketch of a rational reconstruction of Wittgenstein's aforementioned charge that epistemologists "sublimate" the ordinary concept of knowledge. Let us distinguish between two visions of what the epistemologist is trying to achieve concerning the concept of knowledge: the "Received View" and the "Family-Resemblance View". The starting point of both views is linguistic data: a plethora of different uses of the word "to know" in various contexts and linguistic constructions. These uses fall into more or less distinct patterns; and the patterns are more or less similar to one another.

On the Received View, the epistemologist approaches this linguistic data with the aim of constructing a *semantic theory* intended to capture what "to know" means *in its central uses*, or what knowledge *really*, in its core, amounts to. It is striking that there is very little discussion amongst epistemologists as to which uses are the central ones – usually it is simply taken for granted that the philosophical tradition has picked out these core uses a long time ago, and that one can here simply rely on the wisdom of previous generations. There is also little debate over the kinds of theoretical virtues that should guide the construction of the semantic theory. When pressed, most epistemologists would probably say that their theorizing is constrained by the same general virtues that are central in natural scientific work: simplicity, scope, accuracy, and so on. No doubt, simplicity dominates the others: epistemologists like their semantic theories to boil down to a simple formula of the sort "knowledge is justified true belief" or "knowledge is factive", etc. But simplicity comes at a price: it forces the epistemologist to narrow his vision and to leave out of the scope of his semantic theory a substantive number of patterns of use of "to know".

Needless to say, epistemologists working within the paradigm of the Received View recognize that their semantic theory – the theory that is supposed to tell us what knowledge *really* is – does not cover all uses of "to know". And they do have an answer to this problem: the semantic theory needs to be complemented by a *pragmatic theory* and by an *error theory*. The pragmatic theory explains features of the

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

concept of knowledge that cannot be explained on the basis of the semantic theory. For instance, and to stick to our already familiar example, it renders understandable why it seems odd to declare oneself a knower of common-sense certainties. Gricean pragmatic principles do the work. Uses that still remain unaccounted for after the semantic and the pragmatic theories have done their work – non-factive uses of "to know" for example – are addressed by the *error theory*. It highlights the conflations underlying such uses.

One author who is refreshingly straight about the *Received View* is Jonathan Kvanvig (2003). Kvanvig distinguishes between those uses of "knowledge" that are "central to epistemological inquiry, related as it is to theoretical concerns" and those uses that are not. The border line he has in mind coincides with the divide between factive and non-factive uses of "knowledge". As Kvanvig has it, utterances of the form "it used to be known that the earth is flat" are instances of "misspeaking" typical of "undergraduates" at the beginning of their philosophical education (xi, 190); expressions of the form "I just *knew* I was going to fail" (when the speaker in fact succeeded) conflate knowledge and psychological certainty; and locutions involving the "current state of scientific knowledge", or "talk about the present state of knowledge (even though we know that some of what falls under that rubric is false)" are merely "honorific" (2003, 201; 2009, 342). Kvanvig goes so far as to say that "if I were to offer a theory of knowledge, I would not expect it to answer to locutions involving the current state of scientific knowledge". Kvanvig takes such honorific uses to belong not to the semantics but to the "pragmatic dimension of epistemic terminology" (2009, 342).

THE FAMILY-RESEMBLANCE VIEW

The Wittgenstein of my rational reconstruction rejects the Received View. He objects to all three elements.

First of all, Wittgenstein finds indefensible the practice of simply trusting that the philosophical tradition is right about the central cases of "to know". He also insists that the selection of central uses is grossly underdetermined by the linguistic data and the vague and hand-waving references to reflective equilibria or theoretical virtues.

Second, Wittgenstein is deeply suspicious of the semantics-pragmatics divide. His reasons have been spelled out in a string of important papers and books by Charles Travis, and I will not attempt to cover this complex territory in a couple of sentences (Travis 1991, 1997, 2006). But a simpler point concerning the semantics-pragmatics divide can be made even within the limited space available. Since the selection of central uses is underdetermined by the data, so is the decision which features of the data should be captured by the semantic theory, and which features of the data should be accounted for by the pragmatic theory. Should we treat factive uses as the semantic core and non-factive uses as explained by pragmatics (Holton 1997) or should we proceed the other way around (Hazlett 2010)? In the absence of

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

proper theoretical constraints, the decisions to go either way appear arbitrary and unmotivated.

Third and finally, the exclusion of certain uses as "errors" or "merely metaphorical" is – from Wittgenstein's perspective – little more than an arbitrary legislating. Native speakers do not regard non-factive uses of "to know" as confusions or mistakes; why then should the philosopher so look upon them? Was the greatest poet of the German language, Goethe, confused about the meaning of "wissen" when he used it in non-factive constructions? (1969, §8)

Wittgenstein's alternative to the Received View is to take *knowledge* as a family-resemblance concept. On this assumption we can do – and need do – *no better* than distinguish the different patterns in our uses of "to know" and describe their similarities and differences. The proponent of the Received View must find this a disappointing project: Wittgenstein leaves no space for telling us what knowledge really, in its core, amounts to, what "to know" really means. Of course, if Wittgenstein is right, then the expectation of finding a core to knowledge is misplaced from the start.

CRAIG'S GENEALOGY AND THE TWO VIEWS

How does Craig's genealogy relate to the two views? To answer this question we need to reflect on a comment made by Craig in 2007, in a paper that reflects in retrospect on the rationale underlying his 1990 book. He writes that

> … the circumstances that favour the formation of the concept of [proto-]knowledge still exist … otherwise I would have had … no support for my thesis that the method reveals the core of the concept as it is to be found now." (2007, 191)

Craig here makes two points: that we still are, at least sometimes, in the epistemic state of nature, and that *protoknowledge* is the *core* of *knowledge*.

The idiom of "core" and "periphery" is characteristic of the Received View. Craig's use of it reveals something like an "avalanche model" of the genealogy of *knowledge*. The starting point is the concept of *protoknower* qua good informant. As this small conceptual "stone" rolls down the snowy mountain, it gathers around it layer after layer of further conceptual features. (Recall that the subtitle of Craig's book is "An Essay in Conceptual Synthesis".) Once the avalanche ends up in our valley, the core has become almost invisible – but it is still there.

Craig's commitment to the Received View seems to me to be the greatest weakness of his book. Moreover, the commitment to this view is not essential to the very idea of a genealogy. Genealogy is important because it identifies the *protoknowledge* pattern of use as the natural origin of the development of our concept of knowledge: after all, the social situation in which this pattern has its home is simpler and less complex than those of other uses. But this does not justify conceptualizing *protoknowledge* as the core of *knowledge*. We do not think of *homo*

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

erectus as the core or essence of *homo sapiens* just because *homo sapiens* developed out of *homo erectus*.

Indeed, so far from supporting the Received View, the genealogical-*ecological* story – to return to my initial theme – lends support to the idea of *knowledge* as a family-resemblance concept. It does so in two ways. On the one hand, the genealogical-ecological story makes the idea of *knowledge* as a family-resemblance concept *more plausible*: concepts of *protoknowledge* in all its various stages of objectivisation are solutions to various enduring types of social interaction. And it is the assembly of these concepts of protoknowledge that make up our concept of knowledge. On the other hand, the genealogical-ecological story also gives us a more refined vocabulary to speak about the relationship between the different patterns of use: it allows us to replace Wittgenstein's vague talk of "similarities" with the more precise idiom of "genealogical descent" or perhaps even of "adaptation of conceptual patterns to new social niches". That I am here likening genealogy once more to ecology is not arbitrary: the genealogical work in this context comes close to the work of historical linguists – like Williams Croft (2000) – who use ecological-evolutionary models to explain language change. The ecological language comes naturally here also because the natural metaphor of a historicised family-resemblance concept is the phylogenetic tree: the concept of the good informant was the starting point, and there are contexts in which it still survives. But as new social situations have developed and as society has become more complex, new uses of "to know" and its cognates have evolved. There is *no core*, only a (hypothetical) historical beginning.

WHY SOMETIMES COMMON-SENSE CERTAINTIES CANNOT BE KNOWN

The reasoning over the last three sections naturally suggests the following hypothesis:

> (*H*) There are not only patterns of use of "knowledge" on the basis of which common-sense certainties are knowledge-apt. There are also patterns of use of "knowledge" on the basis of which common-sense certainties are *not* knowledge-apt.

Is *H* true? To show that it is is to (partially) vindicate the Wittgenstein of *On Certainty*. Here are three pieces of evidence:

(1) Wittgenstein himself reminds us that we sometimes say "I know p" in order "to assure people" (1969, §424). Clearly such assurance is pointless when the p is question is a common-sense certainty.
(2) "To know" takes interrogative constructions as complements. The same is not true for "to believe" (and its cognates). Compare e.g.

   (a) I know who she is. She knows where to go. He knows when to come.
   (b) I believe who she is. She believes where to go. He believes when to come.

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

The sentences in (a) are all well formed, those in (b) are not. This observation has prompted the comment that "we use the notion of knowledge in describing our aims when we address questions to other people" (Welbourne 1993, 9; cf. Vendler 1979; Kusch 2002). On this pattern of use, common-sense certainties are not knowledge-apt. We do not normally ask questions to which common-sense certainties are the appropriate answers.

(3) Perhaps the strongest piece of evidence is the fact that some epistemologists note that *knowledge* attributions often involve credit, praise, and honoring (e.g. Greco 2003). Of course, it is not correct to say that *all* knowledge attributions come with credit and honor, but a substantive number certainly do. And on those patterns where the link is central, common-sense attributions are not knowledge-apt.

My partial vindication of Wittgenstein's claim according to which common-sense certainties are not knowledge-apt is now complete. There are patterns of use of "to know" on the basis of which it would not be correct to say, "I know that the Earth has existed since long before my birth."

Given how central in my vindication of Wittgenstein is his own work on family-resemblance concepts, the following question will invariably arise. If *knowledge* is a family-resemblance concept – if on some uses common-sense certainties are knowledge-apt, on some uses not – how come Wittgenstein gave such one-sided attention only to latter uses, and so little attention to the former? One quick answer is to register the fact that he did not ignore the knowledge-apt uses altogether:

> Haven't I got it wrong and isn't Moore perfectly right? Haven't I made the elementary mistake of confusing one's thoughts with one's knowledge? Of course I do not think to myself 'The earth already existed for some time before my birth', but do I *know* it any the less? Don't I show that I know it by always drawing its consequences? (1969, §397; cf. §§480, 552)

But this is really only a quick and unsatisfactory reply. The fact remains that *On Certainty* gives far too little attention to the situations and uses in which common-sense certainties are knowledge-apt. Perhaps this is a point where we can do no better than agree that *to understand Wittgenstein is (sometimes) to go beyond him.*

SUMMARY

In this paper I have tried to make plausible the following claims:

(1) There is a space and a need for the project of a social cognitive ecology of our concept of knowledge.
(2) Craig's genealogy deserves a central position in this field. Its strength lies, amongst other things, in its ability to highlight the ways in which conceptual and social change are intertwined.
(3) A critical dialogue between Craig and Wittgenstein is mutually beneficial:

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

(a) Genealogy can provide Wittgenstein with a richer conceptualisation of the similarities between the patterns in family-resemblance concepts.
(b) Genealogy helps to identify the protoknowledge pattern of use on the basis of which Wittgenstein is right to say that common-sense certainties are not knowledge-apt.
(c) Wittgenstein is crucial for pinpointing and correcting Craig's problematic commitment to the Received View of epistemological method.*

REFERENCES

Barnes, B. 2000. *Understanding Agency: Social Theory and Responsible Action*. London: Sage.
Barnes, B., D. Bloor, and J. Henry. 1996. *Scientific Knowledge: A Sociological Analysis*. London: Athlone.
Bloor, D. 1997. *Wittgenstein, Rules and Institutions*. London: Routledge.
Craig, E. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
Craig, E. 2007. "Genealogies and the State of Nature." In A. Thomas (ed.), *Bernard Williams*, pp. 181–200. Cambridge: Cambridge University Press.
Croft, W. 2000. *Explaining Language Change*. Harlow: Longman.
Daston, R. and P. Galison. 2007. *Objectivity*. New York: Zone Books.
Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
Greco, J. 2003. "Knowledge as Credit for True Belief." In M. DePaul and L. Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 111–34. Oxford: Oxford University Press.
Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
Hawthorne, J. 2003. *Knowledge and Lotteries*. Oxford: Oxford University Press.
Hazlett, A. 2010. "The Myth of Factive Verbs." *Philosophy and Phenomenological Research* 80(3): 497–522.
Holton, R. 1997. "Some Telling Examples: A Reply to Tsohatzidis." *Journal of Pragmatics* 28: 625–8.
Kelp, C. 2011. "What is the Point of 'Knowledge' Anyway?" *Episteme* 8(1): 53–66.
Kusch, M. 2002. *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: Oxford University Press.
Kusch, M. 2009. "Testimony and the Value of Knowledge." In A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value*, pp. 60–94. Oxford: Oxford University Press.
Kusch, M. 2011. "Genealogy, Metrology, Historical Epistemology." Forthcoming.
Kvanvig, J. L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
Kvanvig, J. L. 2009. "Responses to Critics." In A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value*, pp. 339–51. Oxford: Oxford University Press.
Moore, G. E. 1925/1959. "A Defence of Common Sense." In *Philosophical Papers*, pp. 32–59. London: George Allen & Unwin.

https://doi.org/10.3366/epi.2011.0003 Published online by Cambridge University Press

Moore, G. E. 1939/1959. "Proof of an External World." In *Philosophical Papers*, pp. 127–50. London: George Allen & Unwin.
Schelling, T. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
Travis, C. 1991. "Annals of Analysis." Critical Notice of H. P. Grice, *Studies in the Way of Words*. *Mind* 100: 237–64.
Travis, C. 1997. "Pragmatics." In B. Hale and C. Wright (eds.), *A Companion to the Philosophy of Language*, pp. 87–107. Oxford: Blackwell.
Travis, C. 2006. *Thought's Footing: A Theme in Wittgenstein's Philosophical Investigations*. Oxford: Clarendon Press.
Vendler, Z. 1979. "Telling the Facts." In P. A. French, T. E. Uehling, and H. K. Wettstein (eds.), *Contemporary Perspectives in the Philosophy of Language, Midwest Studies in Philosophy*, vol. 2, pp. 220–32. Minneapolis: University of Minnesota Press.
Welbourne, M. 1993. *The Community of Knowledge*. Aldershot: Gregg Revivals.
Williams, B. 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press.
Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
Winch, P. 1958. *The Idea of a Social Science and Its Relation to Philosophy*. London: Routledge & Kegan Paul.
Wittgenstein, L. 1953. *Philosophical Investigations*. Oxford: Blackwell.
Wittgenstein, L. 1969. *On Certainty*. G. E. M. Anscombe and G. H. von Wright (eds.). Oxford: Blackwell.
Wittgenstein, L. 1978. *Remarks on the Foundations of Mathematics*. 3rd ed. Oxford: Blackwell.
Wittgenstein, L. 2000, *Nachlass: The Bergen Electronic Edition*. CD-ROM. Oxford: Oxford University Press.

NOTES

* For comments and encouragement I am greatly indebted to Mikkel Gerken, Alvin Goldman, Sarah Gore Cortes, Veli Mitova, Christian Piller, Duncan Pritchard, and the audiences at the *Episteme* conference in Edinburgh and the Philosophy of Science Workshop at the University of Vienna.

Martin Kusch holds a chair in theory of science and epistemology in the department of philosophy at the University of Vienna. Before moving to Vienna in 2009 he was professor of philosophy and sociology of science at the department of history and philosophy of science at Cambridge University. His most recent books are *Knowledge by Agreement* (2002) and *A Sceptical Guide to Meaning and Rules: Defending Kripke's Wittgenstein* (2006). He is currently writing a book on Wittgenstein's challenge to contemporary epistemology.