

Teleologies and the Methodology of Epistemology

GEORGI GARDINER

Rutgers University

Chapter one of *Epistemic Evaluation: Purposeful Epistemology*, eds. J. Greco and D. Henderson. Oxford: Oxford University Press, 2015, pp. 31–45.

Abstract

The teleological approach to an epistemic concept investigates it by asking questions such as ‘what is the *purpose* of the concept?’, ‘What *role* has it played in the past?’, or ‘If we imagine a society without the concept, why would they feel the *need* to invent it?’ The idea behind the teleological approach is that examining the function of the concept illuminates the contours of the concept itself. This approach is a relatively new development in epistemology, and as yet there are few works examining it.

This paper aims to fill this gap and engender further understanding of the teleological method. I first contrast the teleological method with more orthodox approaches in epistemology. I then draw a three-way taxonomy of different kinds of teleological approach and provide an example of each kind. The teleological approach is often presented as antithetical to the more orthodox approaches in epistemology, and so in competition with them. I demur. I argue that the methods can be fruitfully combined in epistemological theorising; in the final section I suggest specific ways the teleological approach can be incorporated alongside more orthodox methods in a general methodological reflective equilibrium.

Introduction

A teleological approach to an epistemic concept investigates it by asking questions such as ‘what is the purpose of the concept?’, ‘What role has it played in the past?’, or ‘If we imagine a society without the concept, why would they feel the need to invent it?’ The idea behind the teleological approach is that examining the function of the concept illuminates the contours of the concept itself. This approach is a relatively new development in epistemology, largely introduced by Edward Craig’s *Knowledge and the State of Nature*; as yet there are few works examining the method.¹

I aim to fill this gap and engender greater understanding of the method. In section one I sketch some orthodox methods in epistemology and explain how the teleological approach differs from these. In section two I examine the teleological approach, draw a taxonomy of different kinds of teleologies and provide an example of each kind. Some theorists suggest the teleological approach is antithetical to orthodox approaches. I argue, by contrast, the teleological approach can be fruitfully incorporated alongside the orthodox approaches within a general methodological reflective equilibrium. I argue for this claim in two ways. Firstly, for each example of a teleology in section two, I draw out what the account suggests about the nature of knowledge, thereby illustrating how the teleological approach informs and complements standard epistemological projects. Secondly, in section three I explain specific ways that insights from the teleological method can augment the orthodox methods. Whilst I focus on understanding epistemic concepts, such as knowledge, much of what I say can be fruitfully generalised to other areas of philosophy.

¹ Klemens Kappel has written a useful and insightful paper on teleologies in epistemology, see Kappel (2010). Kusch (2009) was also helpful in constructing my taxonomy. See also Craig (2007), Fricker (1998: esp. section three) for penetrating discussion about the method.

1. Orthodox Approaches

In this section I sketch three orthodox approaches that can be usefully contrasted with the teleological approach. The descriptions are best viewed as summaries, rather than as thorough accounts. It should also be noted this list is not intended as an exhaustive survey of methods in epistemology, nor do I suggest the methods are mutually exclusive; in practice theorists often use all of the methods to varying degrees, and each generates inputs into a large reflective equilibrium. By teasing the methods apart, however, we are better placed to understand them.

i.) Extension-first approach

The ‘extension-first’ method elucidates a concept using various cases to mine our judgements about the intuitive extension and to test the intuitive extension against a proposed claim or theory. A presupposition of this method is that we have intuitive access to the extension of the concept (that is, we have reliable pre-theoretic judgements about when the proposition *S knows that p* is correct). We test a claim or theory by determining whether the theory’s extension and the intuitive extension align.

Suppose our target theory is *Knowledge is true belief*. We can test this theory by devising a case: Sammy correctly guesses the result of a dice throw. The theory’s extension entails Sammy knows the result; if intuition dictates he does not know then there is a mismatch. According to the extension-first method we must either reject the intuition, and provide a debunking story to explain away the misleading intuition, or we must sanction the intuition, and amend the target theory accordingly, so that it no longer conflicts with the intuitive extension.²

Some judgements about the extension might be harder to explain away than others: The Sammy vignette seems to generate a confident and unambiguous judgement about a relatively central case. Other cases might be more peripheral or obscure, and might generate less confident judgements. These latter intuitions will be easier and less costly to explain away. The idea driving the method is that thought experiments reveal our intuitions about the extension of the concept, and after a process of reflective equilibrium—in this case, a two-way reconciliation between intuitive judgements about particular cases and general principles about the nature of knowledge—we can provide an account of the concept.³

Experimental philosophy can be seen as a variation of the extension-first approach, since it employs cases to collect intuitive responses about the extension of the concept by asking whether the case exhibits the target concept. Whereas traditionally those using the extension-first method consult fewer people, and typically consult

² Following Timothy Williamson’s book *The Philosophy of Philosophy* (2007), there has been a recent surge of interest in determining precisely the content of the judgement we have in response to the vignettes when employing the extension-first method. Various proposals have been advanced, most of which suggest the content is some kind of modal claim. Examples include ‘necessarily, anyone who stands to *p* in the relation described by the vignette has a true belief that *p* but does not know that *p*’, ‘it is possible that someone stands to *p* as described by the vignette, and she has a true belief that *p* but not know that *p*’, and ‘counterfactually, were someone to stand to *p* as described in the vignette, then she would have true belief that *p* but not know that *p*’. See Malmgren (2011: esp. 274–281) for discussion of these proposals. Each of the proposals concerns purely the extension of knowledge. Ernest Sosa (forthcoming) advances a competing theory. He suggests that our judgement on hearing a vignette is that, roughly, ‘The subject of the vignette has a true belief, but does not thereby know, that *p*. Since this judgement includes the ‘thereby’ relation, it is not purely extensional, and so strictly speaking falls outwith the extensional approach as outlined here.

³ In using the expression ‘account of a concept’ I hope to remain ecumenical among an analysis of a concept, a Carnapian explication, a synthetic account such as Edward Craig’s, or other similar understandings of the aim of epistemological theorising.

experts, such as fellow philosophers, experimental philosophy consults large numbers of people, often the untrained “folk”, about the intuitive extension.

ii.) Intension-first Approach

The intension-first approach uses intuitive judgements about the intension of the target concept to inform an account. We believe, for example, that whether S knows that p is not typically sensitive to the day of the week.⁴ Plausibly we access this day-of-the-week-invariance platitude directly from our intuitive understanding of the concept, without proceeding via the intuitive extension. Whilst the day-of-the-week datum is not much help in formulating a theory of knowledge, there may be other judgements about the intension that bear more significantly on our analysis. If we discover from reflecting directly on the intuitive intension whether knowledge has a conceptual connection to rational action or blameless assertion, for instance, or whether possessing knowledge is less demanding for children than for adults, then this will significantly inform our theory.

Theorists who begin with a set of platitudes and build an account of knowledge from them employ the intension-first approach. Such platitudes may include the value platitude (knowledge has value), the anti-sceptical platitude (we have some knowledge), the ability platitude (knowledge is, at least in part, to the credit of the knower), and the anti-luck platitude (knowledge enjoys some kind of modal stability). Theorists who take as inputs the normative roles that knowledge plays in our socio-cognitive environment (such as the norm of assertion, the norm of rational action, or the epistemically proper end of inquiry) can also be fruitfully seen as employing the intension-first approach.

iii.) Linguistic Approach

Another method we can usefully contrast here is the linguistic or ‘ordinary-language’ approach: When analysing epistemic concepts, we survey how people *use* epistemic terms and this guides theorising about the concepts. The linguistic approach might proceed with corpus studies, testing felicity judgements, semantic analysis, cross-linguistic studies, or simply observing, using empirical methods, how people use the term ‘knows’ and its cognates. Studying linguistic behaviour to illuminate concepts presupposes that there are deep connections between language and concepts; it is fruitful only if language provides some insight into concepts.

This method is employed in debates about epistemic contextualism, for example, as how people use ‘knows’ in ordinary language is held to support claims about which parameters affect the truth of knowledge attributions. Similarly we can examine what epistemic terms convey in conversations: ‘Apparently Mark attended the party’ conveys to competent language users that the speaker does not possess first-hand evidence, for instance.

The linguistic approach is distinct from extension- and intension-first approaches: Whereas the extension-first approach asks whether the concept intuitively applies to a case, the linguistic approach examines language use itself. Proper language use—including attributions and denials of knowledge—can diverge from judgements about the extension of knowledge. We might assert a knowledge ascription, even whilst denying the corresponding extensional claim. On hearing a piece of gossip we might exclaim “I knew it!”, whilst recognising we did not know. Here language use conveys

⁴ For some propositions, of course, whether S knows p does depend on the day of the week. S cannot know it is Thursday, for example, unless it is Thursday. But it is not typical for knowledge to depend on the day of the week. Thanks to Ernie Sosa for pointing this out.

something other than a literal knowledge attribution.⁵ We might comfort an ill friend by saying “I know you will survive this” whilst recognising we do not know that she will. We might express annoyance at getting a pub quiz question wrong by asserting, “I knew the answer”, whilst recognising we did not. These utterances might be linguistically appropriate, even though strictly speaking false, precisely because linguistic behaviour is distinct from reporting judgements about the extension. Conversely it might be inappropriate to assert a knowledge claim, even if you recognise that knowledge obtains. Perhaps lottery cases are like this: Perhaps it is typically inappropriate to assert “I know I won’t win the lottery”, even though you do know, because the assertion typically implies you possess some non-statistical, inside evidence.

Another, related, difference between the linguistic approach and the extension-first approach is that the former may be particularly useful in determining various different conversational uses of the term ‘knows’, whereas the extension-first approach may be better suited to exploring just one concept of knowledge.⁶

Those employing the linguistic approach may either count all ordinary language uses of ‘knows’ as equally proper, and so theorise about the nature of knowledge from all uses; or they may rule that some cases are proper or central uses, whilst others are figurative, improper or mistaken. Either way, ordinary language use shapes the account of the concept being developed.

2. Teleological Approaches

Teleological approaches aim to illuminate a concept by asking what the *point* of that concept is; what *purpose* it fulfils, what *need* it meets, what *function* it has, or what *role* it characteristically plays. The idea is that by focusing on the needs the epistemic concept fulfils we illuminate the nature of the concept.⁷

Because it relies less on intuitive responses to individual cases, the success of the teleological approach does not depend on our having reliable intuitive access to the intension or extension of the target concept. It also does not require that our ordinary language behaviour track closely the contours of the concept. The approach is distinct from the others because it begins by looking at our socio-cognitive economy, and determining what concepts might be useful, rather than starting with any particular claims about a concept’s contours, such as particular instantiations or intensions.

There has been a recent surge of interest in teleologies in epistemology, and whilst various teleologies have been proposed, the method has not been extensively

⁵ Perhaps in this case “I knew it!” conveys something like, “That information fits well with my previous beliefs and suspicions”, and generates the perlocutionary force of encouraging the gossip to reveal more gossip.

⁶ A fourth difference between the linguistic approach and the extension- and intension-first approaches, is that plausibly expertise plays a different role in data collection. Perhaps consulting experts is particularly helpful when generating results using the extension- and intension-first approaches; plausibly consulting experts provides more accurate and coherent responses. By contrast, plausibly when seeking information about linguistic behaviour any competent—or native—speakers would suffice, and expertise does not provide extra advantage for, and may even hinder, inquiry. If correct this suggests that insofar as theorists aim to determine the structure of a coherent concept we should ask experts, rather than the folk, and use the extension- and intension- first methods. And insofar as we aim to tap into folk concepts we should consult the folk, and the linguistic approach may be particularly well suited for this.

⁷ Teleological approaches are not limited to understanding concepts: they can also be used to understand, for example, artefacts or cultural phenomena. Kappel (2010) illustrates this with a practical explication of the artefact of a car. I focus here on teleologies as employed to illuminate epistemic concepts.

explored.⁸ In this section I construct a taxonomy comprising three variations of the approach. For each variation I explain the approach, and provide an example for illustration. I then draw out what that example teleology, if correct, suggests about the nature of knowledge.

Many theorists who adopt the teleological approach in epistemology couple their application of it with a rejection of the project of giving an analysis of the concept of knowledge in terms of necessary and sufficient conditions, and instead judge that the project of epistemology ought to be elucidation of concepts, constructing a network analysis, or some other synthetic, non-orthodox, aim.⁹ The guiding assumption seems to be that the teleological approach is not well suited to serving the aim of conceptual analysis.¹⁰ Conversely the orthodox methods are seen as appropriate for conceptual analysis, but ill-equipped for a synthetic aim, so that those who reject the orthodox methods thereby see themselves as having reason to eschew the analytic aim in favour of a synthetic aim.¹¹ In a similar spirit, theorists who adopt the aim of rendering knowledge into constituents, such as necessary and sufficient conditions, may doubt the usefulness of the teleological method for their project.¹²

I think these methodological inferences are mistaken. Against the tide of people working in teleological theorising I hold that once we have distinguished clearly the aim, such as analysing knowledge into constituents or generating a synthetic account of knowledge, from the method, namely positing teleologies or using more orthodox methods, we should see that either method can be fruitfully applied for either aim. In particular I argue the teleological method can be used alongside more orthodox methods to generate inputs for our reflective equilibrium as part of the analytic project. I advance this claim in two ways: In this section I give examples of how specific teleologies can support particular claims about the analysis of knowledge. In the next section I articulate several ways the teleological and the more orthodox approaches can be fruitfully combined.

⁸ For existent discussions of the methodology, see Craig (2007), Fricker (1998: esp. section three), Kappel (2010), and Kusch (2009).

⁹ In line with Craig's terminology, by an "analytic" aim I mean the project of rendering the concept into constituent parts, such as necessary and sufficient conditions (Craig (1990: 2)). By "synthetic" aim I mean something more broad: any aim of better understanding knowledge, or the concept of knowledge, that does not try to analyse knowledge into constituent parts. At least, this is certainly in line with Craig's use of "analytic". It is slightly unclear what Craig means by "synthesis", other than that he contrasts it with analysis. The term "synthesis" and its cognates only appears once in Craig's monograph—twice if you count the subtitle—and not at all in Craig (1986/1987) or Craig (2007). (This is rather like, if you will permit the comparison, learning that Darwin's *Origin* only contains the term "evolved" once.) Craig's not being wholly explicit about what he means by a synthesis may have contributed to a subsequent conflating of his method (a genealogy) with his aim (a non-analytic account of a concept).

¹⁰ For explicit claims to this effect, see Fricker (1998: esp. section three; 2008: esp. 47–48), Craig (1990: esp. chapter one; 1986/1987), MacBain (2004: 193–196). See also Williams (2002) and Kusch (2009: esp. section three).

¹¹ For claims that adopting a synthetic aim in epistemology should be paired with rejection of the orthodox methods, see Kusch (2011: 17–20) and Fricker (2008: 47–50). A notable exception is the knowledge-first movement, which employs orthodox methods whilst rejecting the aim of rendering knowledge into constituent parts.

¹² There are exceptions, see for example Greco (2012: esp. section four) and Pritchard (2011; 2012). Both harness Craig's teleological account as evidence to support their favoured analyses of knowledge. By using the methods in harmony to increase the inputs into a reflective equilibrium, they exemplify precisely the approach I advocate.

i.) Practical Explication

Practical explications are the most straightforward of the three variations. The method is synchronic (it does not probe the past) and factual (it relies on empirical facts rather than speculative claims). The theorist makes the following two claims:¹³

- i. Given a set of facts about our physical constitution, cognitive powers and environment, and a set of aims and interests, we have a particular conceptual need;¹⁴
- ii. The target concept is what actually and currently meets that need.

The practical explication method runs thus: since the target concept fulfils the identified need, determining the exact nature of the need and discerning what would satisfy that need, should illuminate the concept.

Klemens Kappel offers a practical explication of the concept of knowledge.¹⁵ He notes some propositions are important for practical reasons and we are better off, from a practical point of view, if we have true beliefs about them. We thus have a practical reason to inquire. He notes inquiry costs time and resources, and has no natural stopping point—since there are always further uneliminated error possibilities, we always have a *pro tanto* reason to continue inquiring.¹⁶

The need Kappel identifies is for an inquiry stopper; a predicate that expresses, to ourselves and others, the judgement that we should take the truth of a claim for granted in practical deliberation.¹⁷ In Kappel's words, "The sound mind will be in need of a way of expressing the attitude that inquiry has now taken one far enough, and that one shouldn't worry about remaining as yet uneliminated error-possibilities".¹⁸ He hypothesises the concept 'knowledge' meets this need, and suggests S's knowing *p* entails:

P, and S is in a sufficiently good epistemic position with respect to *p*, such that S ought to take the truth of *p* for granted in her practical and theoretical deliberation. Kappel (2010: 78).

We can apply the thesis that the concept 'knowledge' functions as an inquiry stopper to the project of analysing the nature of knowledge. If the concept of knowledge fulfils this need, as Kappel claims, this suggests pragmatic encroachment about knowledge. Pragmatic encroachment holds that practical factors affect whether the agent's epistemic position with respect to *p* suffices for knowledge (or, to formulate it as a semantic thesis, practical factors affect the truth value of a knowledge attribution in a context).¹⁹ Kappel's teleology suggests pragmatic encroachment because if we care about truth for pragmatic reasons, and we care more about the truth of some propositions than others, then we should invest more time inquiring into those propositions that we have more reason to

¹³ I have adapted this schema from Kappel (2010: 72–73). Kappel describes the second claim as one of stipulation—we stipulate that the target concept is what fulfils the need. I think this may be better viewed as a hypothesising, however, rather than stipulating.

¹⁴ As Kappel clarifies, by 'need' he means something that we are better off if we meet, rather than the stricter sense of something without which we cannot function (Kappel (2010: 79–80)).

¹⁵ Chris Kelp (2011) independently posits a similar function—that of an inquiry stopper—for the concept of knowledge.

¹⁶ Kappel (2010: 73–76).

¹⁷ Kappel (2010: 76).

¹⁸ Kappel (2010: 75).

¹⁹ Hannon (2013) develops the view that Craig's genealogy of knowledge, described below, supports pragmatic encroachment about knowledge.

care about. So it takes less inquiry to reach the point of stopping when investigating less important propositions (if, for example, stakes are lower). Thus if knowledge marks when to stop inquiry, it is easier to know less important facts. This is the pragmatic encroachment thesis.²⁰

ii.) *Historical Genealogical Teleology*

I introduced the practical explication teleology, which is synchronic; the other two kinds are genealogical and hence diachronic. They also aim to elucidate a concept by asking teleological questions—examining the function of the concept to illuminate the concept itself—but the genealogical versions are retrospective: they look at what function the concept has played or what need it used to meet. There are two kinds of genealogical teleologies, historical and hypothetical.

The historical approach examines the function of the concept in the actual past, and from this infers facts about the present concept. It uses empirical data from historians about the role the concept played, or the way a term was used, and then analyses the conceptual need that was being fulfilled. The theorist then postulates cultural change that causes the concept to evolve over time and thereby illuminates the contemporary concept. The historical approach postulates a hereditary relationship between ancestor concept and its contemporary counterpart.

Martin Kusch provides an example of a historical genealogical teleological account.²¹ He draws on Steven Shapin's (1994) *A Social History of Truth: Civility and Science in Seventeenth Century England*, which tracks epistemic language and concepts during the Enlightenment. Shapin notes that natural philosophers of seventeenth-century England accepted testimony was needed in order to learn about the natural world, but recognised that finding good testifiers was a difficult practical problem. Seventeenth-century literature suggests several maxims, such as 'assent to testimony which is plausible', 'assent to testimony which is multiple', and 'assent to testimony which is consistent'. As soon as a maxim was suggested, however, problems were found. Only one maxim was never challenged: 'assent to testimony from sources of acknowledged integrity and disinterestedness'.²² But how, asks Kusch, were they supposed to identify who the disinterested reporters were? This is the conceptual need Kusch identifies: finding and predicating those sources who were disinterested and so would tell the truth.

Kusch claims this is where nobility and freedom become important in the genealogy of knowledge. He argues seventeenth century intellectuals believed "Gentlemen were truth-tellers because nothing could work upon them that would induce them to be otherwise."²³ They were believed to be disinterested because of financial independence, and so were considered 'knowers', unlike labourers and females. Kusch quotes Shapin on the servants who assisted Boyle in his lab:

Whatever information the domestics produced, it became knowledge, and thus a property of the gentlemanly community of natural philosophers, only once it was

²⁰ Moves are available, of course, to endorse Kappel's teleology yet deny pragmatic encroachment. Perhaps one could hold that the concept of knowledge serves to flag when we can permissibly cease inquiry in some default, typical practical context. Sometimes, depending on the context, we must inquire more—or may inquire less—than that required for knowledge. Such views might successfully deny pragmatic encroachment whilst endorsing Kappel's view. But Kappel's teleology, at the very least, *prima facie* suggests, and gives evidence for, pragmatic encroachment. Thanks to Lisa Miracchi for pressing me on this point.

²¹ Kusch (2009).

²² Shapin (1994: 212); Kusch (2009: 86).

²³ Shapin (1994: 84); Kusch (2009: 87).

vouched for by Boyle or another gentleman. (Shapin (1994: Ch.8), quoted in Kusch (2009: 87)).

Gentlemen held that employees and women did not have the financial freedom to create and pass on knowledge. Kusch quotes Shapin:

The conventions of gentlemanly experimental philosophy did not allow for anyone to openly express disbelief in a report coming from a gentleman. The situation was very different for all those who did not make the gentry grade: women, servants, “the poor and the mean in general”, merchants, Catholics, Continental gentry, Italians and politicians. In the cases of all of these groups, their “unreliable truthfulness ... was pervasively referred to their constrained circumstances”. (Shapin (1994: 86); Kusch (2009: 87)).

According to contemporary usage these people could possess true beliefs, but not knowledge. The term ‘knower’ was reserved for someone you could trust because they would not need the financial gain made possible by lying.

What does this teleology suggest about the concept of knowledge? Kusch argues his historical genealogy is evidence that “to attribute knowledge is to attribute honour, freedom and social power”.²⁴ He argues that we should expect conceptual connections between the concept of knowledge and that of value, praising someone, virtue, knowledge being an honourific, and bestowing status (such as freedom from being doubted).²⁵

iii.) Hypothetical Genealogical Teleology

The third kind of teleological approach is a hypothetical genealogical teleology. This approach is genealogical as explores the beginning and history of the concept, and hypothetical because rather than tracking the actual history of the concept, it posits a fictional historical narrative.²⁶

The method runs thus: First, imagine a state of nature society, and hypothesise our needs in that state of nature, then posit a proto-concept which fulfils those state-of-nature needs. We then hypothesise how the proto-concept would evolve as a response to societal change, thus illuminating our contemporary concept.

Craig’s seminal *Knowledge and the State of Nature* employs this hypothetical genealogical methodology, and it paved the way for other similar approaches in epistemology, such as Kusch’s and Kappel’s work described above. Here I provide only a quick summary of his teleology. Craig imagines a primitive society in which people are language-using and minimally cooperative. He posits that in this state of nature humans need true beliefs about their environment. Given that inquiry is costly, and some people have better access to true beliefs than others, we have a need to tag good informants.

²⁴ Kusch (2009: 83–87).

²⁵ Kusch has independent reason, inspired by Williams (2002), for positing that knowledge ascriptions are honourific (Kusch (2009: sections six to eight)). Following Williams, Kusch notes that in a state of nature people would free ride on the institution of testimony (Williams (2002: 88); Kusch (2009: 74)). He thus identifies a need: motivating people to investigate and testify honestly, and incentivising people to not free ride on the (state of nature) institution of sharing epistemic goods within a community. He suggests that attributions of knowledge fulfil that need, and—here Kusch departs from Williams—they can do so because knowledge attributions are honourific; they bestow value and praise the subject (Kusch (2009: 79)).

²⁶ Fricker (2008: 47) refers to this as a narrative of ‘semi-fictional time’. Craig (2007: esp. 190–197) discusses to what extent his genealogy, and state of nature genealogies in general, qualify as fictional as opposed to real.

Craig's hypothetical genealogy claims the proto-concept of knowledge fulfils the need of tagging good informants. He then describes a process of increasing objectivity through which the concept changes over time to become less relativised to particular inquirers, so that knowledge attributions can be usefully passed among individuals. 'S knows that p' changes from expressing *S is a good informant for me about whether p* to expressing *S is a good informant for anyone about whether p*.

We can then ask what Craig's account suggests about knowledge. Several suggestions have been advanced; here I focus on one. Duncan Pritchard argues that the concept 'good informant' has two aspects: someone on whom we can rely, and someone who is reliable.²⁷ Pritchard argues that these correspond to modal and virtue-theoretic conditions respectively. Thus we should expect an analysis of knowledge to reveal a bipartite structure, reflecting these dual aspects of 'good informant'. Thus, he concludes, Craig's thesis that knowledge attributions tag good informants provides reason to think the modern concept 'knowledge' has the structure of anti-luck virtue epistemology (the thesis that knowledge is safe, virtuous, true belief).

3. Roles for Teleologies

I have introduced the teleological approach, explained three variations of it, and contrasted the approach with more orthodox methods in epistemology. In this section I explain specific ways teleological theorising can complement the orthodox methods. In doing so I hope to further advance the view that teleological theorising can augment, rather than compete with, the orthodox methods.

Firstly, deploying the teleological approach may reveal how precise or fine-grained we should expect a concept to be, and so delineate the limits of the other three methods. For example, the concept 'adult' serves a legal need (voting, criminal law etc.). It requires a relatively precise definition to fulfil that need—and sure enough it has one. In contrast the concept 'adolescent' fulfils needs in marketing, media, social science, and stereotyping. In order to fulfil these roles the concept does not need such precise contours; its boundaries can be more vague. Thus we have reason to think it is misguided to seek a precise analysis of the concept 'adolescent' and we ought only aim for more broad-brush analyses.

To illustrate within epistemology: those working on anti-luck epistemology employ myriad subtle variations of examples, such as variants of barn façade cases, to mine intuitions about epistemic luck's exact nature.²⁸ These cases differ only slightly from one another; the distinctions the theorists adjudicate among are fine-grained. Perhaps exploring the function of the anti-luck aspect of knowledge will indicate that we should not expect the anti-luck condition to be quite so fine-grained, and our intuitions about the extension to be quite so discriminatory. Thus the teleological method could illuminate the limits of the other methods.

Secondly, teleological accounts engender understanding by making intelligible why we have the epistemic concepts that we have. We can discern the structure of knowledge using orthodox methods, such as the extension-first approach, and then employ teleologies to explain this structure. The thesis that knowledge functions as an

²⁷ Pritchard (2011; 2012).

²⁸ Examples might include Kripke's red barn cases (Kripke (2011: chapter seven)), cases where there are forks in the road near barn façade county and only some paths would lead the protagonist to barn façade county, and cases where the façades were removed the previous day, leaving only real barns on the day S drives through. Some theorists might already worry about whether we can adjudicate whether S knows in the simple barn façade cases, and as the variations become increasingly baroque, these concerns only heighten.

inquiry stopper, for instance, can help explain the role of truth in knowledge. Teleologies deepen understanding by explaining how constituents of the concept come together to perform its function.²⁹ As Craig points out the teleological method can explain features of knowledge, such as why the term can be found in every language, and can cast light on the appeal of various proposed analyses of knowledge.³⁰

Thirdly teleologies can provide independent verification of the results of other approaches. The teleological approach is less dependent on ‘raw’ intuition, and not in the direct way of the intension- and extension-first approaches. Comparing results from the teleological approach might indicate that our raw intuitions about cases are reliable. This can happen if, for instance, the intuitive extension comports with a proposed teleology. Again, these methods interact in reflective equilibrium, if no plausible teleology fits an account developed using the orthodox methods, then this is a strike against the account’s plausibility.

Conversely, results from orthodox methods can provide justification for, or adjudicate among, given teleologies. There is nothing bedrock about a chosen teleology or what we claim it reveals about the concept: we judge the plausibility of a teleology by determining whether the extension or intension predicted by the teleology matches the intuitive one.³¹ Someone committed to the falsity of pragmatic encroachment may well reject Kappel’s teleology, for example. And they will have developed their opinion on pragmatic encroachment via the orthodox methods of intension-first, extension-first and linguistic studies. Whilst few theorists hold that intuitions about the extension and intension of a concept are infallible, or even enjoy the level of confidence that observations typically have in empirical science, they may be closer to justificatory bedrock than theses about the function of the concept, and so can justify proposed teleological accounts. In general teleological studies increase the kinds of sources that feature in our reflective equilibrium, which increases the justification of the resultant theory.

Finally, the teleological approach might play a role in responding to sceptical challenges. The teleological approach is non-sceptical in a flatfooted way. There are few candidate roles for the concept of knowledge that entail the extension of the concept is empty. There is the role of picking out the mental state, evidence base, or epistemic

²⁹ It should be noted, however that in some special cases the teleological approach cannot illuminate why we have the concepts we do. Perhaps, for example, the concept’s structure came about through mistake, equivocation or misunderstanding, rather than by being well adapted to meet a need. In such cases no functional story will explain the structure.

³⁰ Craig (1986/1987: 212) writes, “Suppose that the problem of the analysis had been solved, so that agreed necessary and sufficient conditions for the ascription of knowledge were now on the table. Many writers make one feel that this would be a terminus, and investigation concluded. I should see this as a prolegomenon to further inquiry: why has a concept demarcated by those conditions enjoyed such widespread use?” Craig sees the traditional project as concerning the word “know”, rather than either the concept of knowledge or knowledge itself (Craig (1986/1987: 221)). Many theorists deny this claim. But the methodological point he makes is important: there is more to theorising about knowledge than simply generating necessary and sufficient conditions, or capturing the intuitive extension with an explicit intension.

³¹ One worry is there is not enough to constrain teleological theorising, especially regarding hypothetical and historical teleologies. Fricker writes, “State of nature stories... are notorious for providing a blank canvas onto which a philosopher may paint the image of his personal theoretical predilections” (Fricker (1998: 164)). Kappel (2010: 80–81) expresses a similar worry. The concern is that if teleological theorising is too open ended, we cannot appropriately adjudicate among proposed teleologies, and teleologies cannot verify or calibrate the results of other methods. I am sympathetic to this worry, but see, for example, Pritchard (2011) for an example of using a given teleology to adjudicate between competing analyses, and see Kelp (2011) for an example of employing extension-first considerations to adjudicate between competing teleologies.

relation to the world that is unobtainable for people in the human epistemic condition. There is flagging an ideal or perfect epistemic position. Perhaps there is the role of describing an epistemic haven that angels and gods can enjoy, but lesser beings could never enter. These kinds of functions for the concept of knowledge would result in a sceptical conclusion: we do not possess much, if any, knowledge. But these roles for the concept of knowledge are not compelling; our behaviour and thought indicates the concept does not fulfil these roles. Instead it seems the concept of knowledge fulfils roles that entail a broad extension; we possess much of the knowledge that we typically take ourselves to have. In this way the teleological approach is a non-sceptical approach to epistemological theorising.

But the teleological approach can provide resources to engage with scepticism in less flatfooted ways too.³² One fruitful way to understand the sceptical challenge is as an attempt to deny us something that we thought we possessed, and that we care about possessing. Perhaps there are some epistemic states, practices, or competences—such as Cartesian certainty about the external world, perfectly trustworthy testimony, or wholly infallible thinking, for example—that sceptical reasoning shows we do not possess. If we either do not value those things, or on reflection would recognise that we do not have them, then arguably relinquishing these things to scepticism is not so troubling or serious. By contrast there might be other states, practices, and competences that we do value, and that we take ourselves to ordinarily possess. Examples might include the legitimacy of our practices of giving and accepting reasons for belief, typically being in a position to testify responsibly, and typically being warranted in trusting our reasoning, perceptual and memorial faculties. Relinquishing these things to scepticism would be a more serious defeat.

The teleological approach can both help explain why we care about the epistemic states, practices and competences that we value, and provide evidence that these phenomena obtain. This is because the approach can illuminate the roles these phenomena play in our social, cognitive, and practical lives, and so help explain their value whilst proving their existence. Perhaps, for example, teleological reasoning illuminates why the institution of testimony has value, and by subsequently showing that the sustainability of this institution requires the concept of knowledge to incorporate honourific force, we thereby illuminate the value of knowledge.³³ To put the point slightly differently: scepticism should show both that something is lacking, and that the thing missing has value. The teleological approach can both provide evidence that something obtains—if the concept of knowledge plays the role of inquiry stopper, for example, this suggests that there is some knowledge—and it can illuminate why that thing has value—because, for example, appropriate cessation of inquiry plays an indispensable or constitutive role in our cognitive and practical lives.³⁴

³² Thanks to Yuval Avnur, Jon Garthoff, and John Greco for helpful discussion of these issues.

³³ See Williams (2002), Kusch (2009: 77ff).

³⁴ Williams (2002) and Nietzsche both posit genealogies to explain features of normative, value-laden phenomena. Nietzsche explores a genealogy for morality; Williams explores one for epistemic value. They both posit that the original function of the value associated with these domains is to manipulate people. Attributions of epistemic value, according to Williams, motivate people to not free ride on the practices of information sharing in the state of nature (Williams (2002: 88–93)). Attributions of moral value, according to Nietzsche, are to motivate the powerful to consider the interests of the less powerful. But whereas Nietzsche's genealogy purports to debunk or undermine the value in question, Williams's genealogy seeks to vindicate it.

Conclusion

My aim is to engender understanding about teleological approaches in epistemology. I outlined the teleological method and distinguished it from the extension-first, intension-first and linguistic approaches. I constructed a three-way taxonomy of the teleological approach and provided an example of each. I illustrated how the teleological approach can be fruitfully applied to the more orthodox aim of analysing the concept of knowledge by drawing out what particular teleologies suggest about knowledge. Contrary to those who believe the teleological approach is in opposition to more orthodox methods, I develop the idea that teleologies can be fruitfully employed alongside orthodox approaches when analysing knowledge. I do this, in part, by explaining some specific ways to do so.

Acknowledgements

Thanks to Yuval Avnur, Peter Baumann, Matthew Chrisman, Megan Feeney, Will Fleischer, Miranda Fricker, Jon Garthoff, John Greco, David Henderson, Chris Kelp, Martin Kusch, Michael Lynch, Peter Milne, Lisa Miracchi, Duncan Pritchard, Ernie Sosa, Chris Tweedt, and Tedla Woldeyohannes for helpful comments and discussions. Thanks also to audiences at the *Episteme* Conference on Cognitive Ecology, St. Louis University, the University of Edinburgh, Swarthmore College, Long Island Philosophical Society, and the Pacific Division Meeting of the American Philosophical Association.

Bibliography

- Craig, Edward (2007). 'Genealogies and the State of Nature' in A. Thomas (ed.), *Bernard Williams*. Cambridge University Press, pp. 181–200.
- _____. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.
- _____. (1986/1987). 'The Practical Explication of Knowledge', *Proceedings of the Aristotelian Society*, Vol. 87: 211–226.
- Fricker, Miranda (2008). 'Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time', *Philosophical Papers*, 37 (1): 27–50.
- _____. (1998). 'Rational Authority and Social Power: Towards a Truly Social Epistemology', *Proceedings of the Aristotelian Society*, 98: 159–177.
- Greco, John (2012). 'A (Different) Virtue Epistemology', *Philosophy and Phenomenological Research*, 85 (1): 1–26,
- Hannon, Michael (2013). 'The Practical Origins of Epistemic Contextualism', *Erkenntnis*, 78: 899–919.
- Kappel, Klemens (2010). 'On Saying that Someone Knows: Themes from Craig' in A. Haddock, A. Millar, & D. Pritchard (eds.), *Social Epistemology*. Oxford University Press, pp. 69–88.
- Kelp, Chris (2011). 'What's the Point of "Knowledge" Anyway?', *Episteme*, 8: 53–66.
- Kripke, Saul (2011). *Philosophical Troubles: Collected Papers*, Vol. 1. Oxford University Press.
- Kusch, Martin (2011). 'Knowledge and Certainties in the Epistemic State of Nature', *Episteme*, 8: 6–23.
- _____. (2009). 'Testimony and the Value of Knowledge' in A. Haddock, A. Millar, & D. Pritchard (eds.), *Epistemic Value*. Oxford University Press, pp. 60–94.
- MacBain, James (2004). 'Epistemic Analysis and the Possibility of Good Informants', *Principia*, 8 (2): 193–211

- Malmgren, Anna-Sara (2011). 'Rationalism and the Content of Intuitive Judgments', *Mind*, 120 (478): 263–327.
- Nietzsche, Friedrich. *On the Genealogy of Morals and Ecce Homo*, trans. Walter Kaufmann, New York: Vintage Books, 1967.
- Pritchard, Duncan (2012). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy*, 109: 247–279
- . (2011). 'The Genealogy of the Concept of Knowledge and Anti-Luck Virtue Epistemology' in S. Tolksdorf (ed.), *Der Begriff des Wissens/The Concept of Knowledge*, de Gruyter, pp. 159–178.
- Shapin, Steven (1994). *A Social History of Truth: Civility and Science in Seventeenth-Century England*. Chicago University Press.
- Sosa, Ernest (forthcoming). 'The Metaphysical Gettier Problem and the X-Phi Critique' in R. Borges, C. de Almeida, and P. Klein (eds.), *Explaining Knowledge: New Essays on the Gettier Problem*. Oxford University Press.
- Strawson, P.F. (1992). Chapter Two 'Reduction or Connection? Basic Concepts' in *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford University Press.
- Williams, Bernard (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton University Press.
- Williamson, Timothy (2007). *The Philosophy of Philosophy*. Blackwell.