

This paper has been published in *Metaepistemology: Realism and Antirealism*. Palgrave Macmillan [2018].

The Genealogy of Relativism and Absolutism

Martin Kusch, University of Vienna & Robin McKenna, University of Vienna¹

Abstract

This paper applies Edward Craig's and Bernard Williams' 'genealogical' method to the debate between relativism and its opponents in epistemology and in the philosophy of language. We explain how the central function of knowledge attributions -- to 'flag good informants' -- explains the intuitions behind five different positions (two forms of relativism, absolutism, contextualism, and invariantism). We also investigate the question whether genealogy is neutral in the controversy over relativism. We conclude that it is not: genealogy is most naturally taken to favour an anti-realism about epistemic norms. And anti-realism threatens absolutism.

Keywords

Edward Craig; Bernard Williams; Paul Boghossian; epistemic relativism; semantic relativism; epistemic antirealism; invariantism; contextualism

Introduction

In this paper we bring together two strands of epistemological theorizing that have lived separate lives: the controversy around epistemic relativism (e.g. Boghossian 2006) and the debate over 'the genealogy of knowledge'. The latter refers to a method first suggested in Edward Craig's *Knowledge and the State of Nature* (1990), and then further developed in Bernard Williams' *Truth and Truthfulness* (2002) and Steven Reynolds' *Knowledge as Acceptable Testimony* (2017). There are different ways to interpret the genealogical method. On our reading (Kusch and McKenna, forthcoming), genealogy comes into its own when used to explain intuitions underlying different epistemological theories. The explanation in question is functional: the *explanans* are ubiquitous human needs; and the primary *explananda* are social practices and institutions of gathering and sharing information about the natural and social world. Specific epistemic concepts such as *knowledge*, and the intuitions involved in their usage, are theorized as arising in this context.

¹ Authors listed in alphabetical order.

Recent years have seen a flurry of publications for and against different forms of relativism. As far as epistemic relativism is concerned, the most influential discussion is Paul Boghossian's *Fear of Knowledge* (2006). Boghossian attacks what he regards as unacceptable forms of epistemic relativism in Richard Rorty (1981) and the 'Sociology of Scientific Knowledge' (=SSK) (e.g. Barnes and Bloor 1982). Boghossian's discussion of relativism foregrounds two ideas: that the relativist denies the possibility of absolute standards of epistemic justification ('non-absolutism'); and that she assumes a plurality of epistemic systems (of epistemic standards). At least in some cases these epistemic systems licence incompatible verdicts on epistemic justification ('pluralism').

A second relativism-debate has emerged over the last fifteen years in philosophical semantics. One key question concerns the semantics of knowledge attributions. 'Contextualists' (e.g. DeRose 2009) hold that the meaning of 'to know' (and its cognates) varies with the context of utterance; 'invariantists' hold that 'to know' has the same meaning in all contexts, and that knowledge attributions are true or false absolutely (e.g. Brown 2006, Rysiew 2001). Relativists (e.g. MacFarlane 2014) hold that, while 'to know' has the same meaning in all contexts, knowledge attributions are true or false only relative to the 'context of assessment'.

We shall bring genealogy to bear on both of these debates, though we focus on the first. We shall call the (anti-)relativism at issue in Boghossian 'B-(anti-)relativism' and the (anti-)relativism at issue in MacFarlane 'M-relativism'. We begin by giving a brief explanation of the genealogical method. We then give genealogical accounts of the intuitions underlying B-(anti-)relativism and M-(anti-)relativism. Up to this point, our treatment of all sides to these disputes will be even-handed: the genealogist seeks to explain the intuitions underlying each view, not to undermine them. In the final section we investigate whether the genealogist can (or should) maintain this neutral stance. Our tentative conclusion is that genealogy is closer to B-relativism than to absolutism.

Genealogy -- A Primer

In *Was Wir Wissen Können* Craig situates his project in close proximity to two unlikely bedfellows: Wittgenstein and natural science (1993: 37). Wittgenstein is an ally since he opposes conceptual analysis in terms of necessary and sufficient conditions, studies the function of concepts, and introduces the category of family-resemblance concepts. Craig's project has affinities with natural science in its method of hypothesis testing, the search for explanation, and focus on evolution. We would add *model-building* to the list: the building of simplified models of complex target systems.

Craig's model-construction has two stages: the first focuses on the 'epistemic state of nature', that is, a small community of language-using humans, engaging primarily in face-to-face communication, who are co-operative, dependent upon one another for information, and of unequal skills and talents. The central question is: Why would a concept like *knowledge* be introduced under these idealised--simplified and distorted--conditions? Craig answers that people in this situation have a salient need, to wit, to pick out and 'flag good informants'. And the concept used to flag good informants is the core--or one central aspect--of *knowledge*.

In the state of nature, individuals depend upon one another for information. Distinguish between the 'inquirer', who needs information that they are currently unable to obtain themselves, and the 'informant', who 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 the concept of a *protoknower* is the central conceptual tool for dealing with this problem; our ancestors used this concept as a tag for good informants. Which conceptual components should *protoknower* contain? Craig's answer is that the concept *protoknower* (*whether p*) comprises these 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 inquirer 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 on (vi), the inquirer needs to find 'indicator-properties' that she can detect and that correlate closely 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 for some inquirers in the state of nature when *p* is the proposition that a tiger is approaching the village. Usually more than one property 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 these elements are present in prototypical situations, the concept has a use even when some elements are missing. Finally, *protoknowledge* differs from *knowledge* in that: (a) only the former is closely tied to testimony; (b) *protoknowledge* is indexed to the capacities and needs of specific inquirers (1990: 90); (c) *protoknowledge* can only be ascribed to others, not to oneself; and (d) protoknowledge is not undermined by accident or luck.

Craig goes to great lengths to show that his model of the epistemic state of nature passes the test of (what the philosophy of scientific models calls) 'external validation'. He does so by arguing that his model predicts and explains several features of our concept(s) of *knowledge* that have been identified in various philosophical theories. These theories are often seen as excluding one another, but Craig thinks that his model can partially vindicate all of them: they contradict each other only if we over-generalise them.

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 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' (1990: 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 with our concept of *knowledge*: 'The concept of knowing ... lies at the objectivised end of the process; we can explain why there is such an end, and why it should be found worth marking in language' (1990: 90-91).

The second stage of Craig's model construction adds a *dynamic* dimension. The dynamic model takes the epistemic state of nature as its starting point and tracks how the concept *knowledge* would evolve and diversify as the simplifications and distortions of the state of nature are removed step by step. This suggests that the dynamic model is really a form of 'de-idealisation'.

The dynamic model too needs to pass muster as far as external validation is concerned. Craig suggests that it correctly predicts, or at least makes sense of,

- contexts with very high epistemic standards (1990: Ch. X),
- intuitions about lottery propositions (1990: XI), and
- our conflicting intuitions about epistemological scepticism (1990: XII-XIII).

As Craig emphasises more clearly in 2007 than in 1990 or 1993, this is not to be taken as a historical thesis: the epistemic state of nature is not a historical period ‘like the Pleistocene’. It is rather a ubiquitous and important type of social-epistemic situation that one is likely to find in all human communities, past and present (2007: 191). This suggests that what the two models present as different stages in the historical development of *knowledge* are really two different types of situation that we experience from day-to-day. In some situations, we are still in the ‘state-of-nature’, in other situations we are at various stages of the process of objectivisation. Note however that this interpretation of Craig shows that his talk of a ‘core’ of *knowledge* might be misleading: if the other uses co-exist, why assume that ‘flagging good informants’ is more fundamental than the other uses? Or put differently, why assume that the right model for the conceptual development is an avalanche rather than a phylogenetic tree? The avalanche model suggests a small conceptual ‘stone’ rolling down the snowy (semantic) mountain, in the process putting on layer after layer of further conceptual features. The phylogenetic tree is without a core. We do not think of *homo erectus* as the core or essence of *homo sapiens* just because *homo sapiens* developed out of *homo erectus*. Of course, to keep with the analogy we here assume counterfactually that *homo erectus* might still be alive today.

Up to this point our exposition of genealogy has focused on Craig. But for certain purposes it is important to complement his account with the developments introduced by Williams (2002) and Reynolds (2017).

From Reynolds we adopt an account of how we collectively impress upon each other the need to testify only what one knows. His guiding thought is that it would be exceedingly laborious for inquirers to keep track of the different degrees of reliability of several informants concerning numerous topics of interest. Clearly, it would be better for inquirers if, prior to volunteering their testimony, informants checked whether their views are appropriately based on the right kind of experience, appropriate norms of epistemic conduct, and good judgment. In short, it would be best if informants internalized the correct ‘testimonial norms’ (2017: 50). Reynolds suggests that we

could collectively bring it about that informants engage in time-consuming epistemic self-monitoring via reward and punishment. When informants tell the truth in relevant ways, we repay them with praise and reciprocation. When informants don't, we sanction them with criticism and lack of cooperation. This practice, Reynolds continues, would eventually lead to the emergence of a special vocabulary 'to express this approval and disapproval and to indicate the conditions that tended to lead to it' (2017: 57). Our concept of *knowledge* is part of this vocabulary.

Williams (2002) adjusts Craig's genealogical method in one crucial respect. According to Williams, Craig's 'imaginary genealogy' needs to be complemented by 'real genealogy', that is, by an engagement with historical and cultural contingent realities. Although Williams does not put it in these terms, one can read him as offering an imaginary and real genealogy of the *social institution* of testimony. Thus Williams shares a method with Craig, but he is more interested in the *values, motivations* and *virtues* that underpin testimony than in *concepts*. According to Williams, the central virtues of testifiers are "accuracy" (a disposition to seek the truth and to report it) and "sincerity".

Unlike Craig, Williams does not take assume that individuals in the state of nature are cooperative and eager to offer information to others. Williams recognises that the institution of testimony is a collective good. Individuals who are rational in a purely self-interested way will try to 'free-ride': they will seek to obtain accurate and sincere testimony from others without offering anything in return. After all, collecting useful information usually involves costly 'investigative investments' (2002: 88).

How is the problem of collective action solved? The core of Williams' solution to the Free-Rider Problem is the suggestion that accuracy and sincerity (and with them the institution itself) must come to be regarded by community members as *shared intrinsic* – rather than as merely *instrumental* – *values* (2002: 90). For community members to have trust in others' reports, they must be convinced that accuracy and sincerity are non-negotiable. And this implies that these values are not – except under extreme circumstances – weighed against, and possibly outweighed by, other interests and values (2002: 91). Williams holds that values come in socially shared webs and systems, that intrinsic values occupy a central position in such webs or systems, and that values make sense only in their essential relations to one another (2002: 92). Williams' view of intrinsic values relates directly to his distinction between imaginary and real genealogies. While no community can exist without values like accuracy and sincerity, different communities embed such values in different 'wider range[s] of values'. And what these wider ranges of values are, 'varies from time to time and culture to culture' (2002: 93).

Williams teaches that the collective good of the testimonial institution can exist only if the virtues of accuracy and sincerity (amongst others) become intrinsic values. But how can communities make sure that their members look upon these virtues as intrinsic values? Williams has an answer to this problem as well. He tells us that ‘people may be discouraged or encouraged, sanctioned, shamed, or rewarded with respect to this’ (2002: 44); that a crucial motivation is the ‘fear ... of disgrace in one’s own eyes, and in the eyes of people whom one respects and who one hopes will respect oneself’ (2002: 116), that ‘the motivations of honour and shame play an important part’ (2002: 120), and that the structure ‘of mutual respect and the capacity for shame in the face of oneself and others, is a traditional, indeed archaic, ethical resource, but it is still very necessary’ (2002: 121). In a nutshell, Williams’ idea is that a system of sanctioning – of honouring and dishonouring – encourages community members to constitute and maintain the institution of testimony and thus the virtues of accuracy and sincerity.

Putting this in different terms, sanctioning in support of the collective good works through the “deference-emotion system” (Scheff 1988). The precondition of this system is our emotional need to continuously monitor how others treat and think of us. We respond to our assessments of this treatment by changing our position on an internal scale that ranges from pride to shame. When we believe that others treat us with deference, when we believe that others honour us, we feel pride (and related feelings) and move ourselves up on the pride-shame scale. When we suspect that such deference and honouring are missing, we tend to feel bad about ourselves and slide downwards towards the shame end of the scale. This emotional dependence on others is exploited by the deference-emotion system. The granting or withholding of deference constitutes a subtle system of social sanctions, a system that we barely notice. And the operation of the deference-emotion system is inseparable from our ongoing conversation about the collective good. Working with and through the deference-emotion system, this conversation continuously re-establishes the importance of the collective good in everyone’s mind.

Finally, we can connect Williams’ and Reynolds’ proposals as follows. Attributions of (proto-)knowledge (and their cognates) play a key role in the collective action that constitutes the institution of testimony. They do so by *honouring* informants. In keeping with the genealogical method, let us first see how the link between protoknowledge attributions and honouring works under the simplified conditions of the state of nature. To publicly apply the concept *protoknower* to someone is not only to classify them as a reliable source of information, it is also to honour them, or to encourage others to do likewise. To classify someone as a protoknower is to praise them for their contribution to the institution of testimony, and thereby for their contribution to the well-being of

the community itself. After all, the community cannot survive without the institution of testimony. By means of attributions of protoknowledge members of the community honour good informants for contributing their part to the existing and flourishing of the community. *Mutatis mutandis*, withholding or denying protoknowledge is a way of censoring and dishonouring. It is to mark someone as not willing, not worthy or not able to participate in the constitution of the collective good, and thus as not fit to be a (working) part of the group. In that sense, to deny someone *protoknowledge* is to expel them from the group.

Our suggestion connects Craig's focus on conceptual needs of the inquirer with Williams' emphasis on the motivational problems of the informants. By using the concept of *protoknowledge* to both classify and honour informants, the inquirer manages to serve two key goals at once: the goal of tagging good informants for future reference (to herself and others), and the goal of motivating community members to make, or keep making, investigative investments.

A Genealogy of B-(Anti-)Relativism

To offer a (Craigean) genealogy of an epistemological view is to identify and explain the intuitions underwriting the view. The explanation tells us why these intuitions are natural given the context of a community of humans operating and maintaining the social institution of testimony. We now turn to the task of offering such explanations for five views: B-absolutism, B-relativism, invariantism, contextualism, and M-relativism. We begin with B-absolutism.

It is helpful to distinguish between B-absolutism and 'monism'. Monism with respect to a domain D is the view that there is only one set or system of norms operating in D. The system might be more or less complex, and contain more or less distinct principles. But monism holds that all these principles cohere with one another, and do not – when interpreted correctly, or applied singly or in combination – contradict one another. The opposite of monism is 'pluralism'. Absolutism takes a crucial step beyond monism: it holds that the one and only system of norms has the property of being uniquely and absolutely correct. Non-absolutist monists refuse to take that step. They acknowledge that there is only one system yet do not draw absolutist conclusions. They might offer an evolutionary explanation why we have the system we have, or they might hold that it is a sort of convention (cf. Hazlett 2014). We are interested in the intuitions that support either one, or both, of these views.

We begin with monism. Here Craig's account of objectivisation is relevant. As objectivization progresses, the standards for useful testimony come to be set in ways that increasingly abstract

from the particular circumstances of inquirers. In the process, testimonial norms concerning correct investigative investments, appropriate levels of accuracy and sincerity, get more demanding. To make sure that the information gained and testified to is of use to an ever wider range of differently situated inquirers, informants must be motivated to meet ever higher demands of reliability. This development is clearly inseparable from the emergence of strongly monistic ways of thinking about testimony. What *every* inquirer, however situated, needs, can no longer vary from one context to another. Objectivisation pushes all parties involved into assuming – at least tacitly and on an intuitive level – that there is but one epistemic system of norms.

We can complement the Craigean genealogy of monism with themes from Reynolds and Williams. Recall the idea that the social institution of testimony is maintained and secured via the deference-emotion system. For this to work, the system of norms used for evaluating epistemic performances had better not to be excessively complicated and context-sensitive. If it were complicated and context-sensitive, it would be beyond the ken of most non-specialists. And, at least under the simplified conditions of the initial genealogical setting, epistemic specialists do not feature in the story. Monistic intuitions – giving support to one epistemic system binding all – are the obvious solution to this problem of evaluation.

This intertwining of epistemic norms and the deference-emotion system can also be used to give a genealogy of absolutism. Adherence to the one and only epistemic system needs to be secured by continually re-establishing its unique value and significance in everyone's mind. Community members are likely to go about this task by giving the values underwriting their epistemic system a central place in their overall web of values. These values therefore come to be regarded as intrinsically valuable, and as beyond all instrumentalist weighing up of different goals. The endpoint is that the epistemic system is itself seen as absolute – as not relative to anything else.

Let us now turn to B-relativism. How might pro-B-relativist intuitions live alongside the absolutist intuitions mentioned above? Again, we will distinguish between two steps: the step towards a plurality of norms, and the further step to denying B-absolutism.

Craig's imaginary genealogy starts with differently situated inquirers focusing exclusively on what is useful to them, taken individually. At this stage, the needs of each individual are the standard against which the utility of information is measured. This clearly is a pluralist stage. On our understanding of Craigean genealogy, this stage is never completely left behind. Even after objectivisation, we still often assess testimony in terms of what we want in our own very specific contexts. In other words, our highly idiosyncratic needs and the idealized abstract needs of the assumed 'very demanding inquirer' exist side by side. Sometimes they overlap, often they do not. Of

course, the two endpoints of the scale – the wholly idiosyncratic and the wholly universal – are both idealizations. Most epistemic evaluations involve the informational needs of different groups of varying size and composition. To be a competent social actor is to have acquired the competence to judge which testimonial norms are relevant for which group. Subject matter, the required degree of reliability, the costs attached to false beliefs, and much else, are relevant variables. The result will be a variety of testimonial norms.

Admittedly, a *variety* of norms is still compatible with monism. We get pluralism only once this variety is partitioned into distinct sets such that different sets give incompatible answers to the same epistemic questions. One system allows one to judge that, in light of Jones' investigative investment concerning *p*, Jones is a good informant as to whether *p*. Another system compels one to say that despite his investigative investment Jones is *not* a good informant in saying *p*. It is plausible to think that the idea of such distinct 'epistemic systems' goes hand in hand with a specific social change: to wit, the forming of distinct 'epistemic communities' with a clear sense of their distinctive shared individual, or collective, goals or needs. There is nothing in Craig's, Reynolds' or Williams' genealogies that makes the emergence of such subgroups intelligible. But we need not add much to the original versions to create the space for this obvious idea. After all every social theory of the development of human societies emphasizes the crucial role of division of labour. And division of epistemic labour is tantamount to a plurality of epistemic communities with distinct standards.

Again, it is helpful to connect this to our collective operation of the deference-emotion system. Consider what might happen when we sanction someone by refusing to accept their testimony, and chastise them for their lack of investigative investment. This will often lead to an aggressive encounter where criticisms fly back and forth. The resulting conflict might endanger social cohesion and peace. The bad informant might well refuse to correct their ways, and be unwilling to share their information with us even on later occasions where their investigative investment is, by our lights, exactly at the required level.

The potential social explosiveness of distrust creates pressure to help the other 'save face'. 'Face saving' might take different forms. We might try to find mitigating circumstances: perhaps the testifier gained their (bad) information in ways we find excusable. Or we might try to find ways of holding that, by their lights, they were justified: perhaps the information was obtained from a usually reliable source, or they were relying on a source they were *prima facie* entitled to trust since they had no evidence suggesting otherwise. Perhaps our testifier acquired a system of epistemic norms different from ours by trusting their elders, which is presumably how we acquired our system

of norms. We therefore have a variety of ways to avoid sanctioning them. We might let things rest and accept that, at least for certain questions, they should not be our first port of call.

Epistemic division of labour and epistemic face-saving incline epistemic agents towards the intuition that there are different epistemic systems at least some of which might give conflicting verdicts on the same issues. The move from this pluralistic stance to the denial of absolutism is of course a further step. It is not obvious that it is inevitable; though once pluralist intuitions are in place, and strong, it may be irresistible.

A Genealogy of M-(Anti)-Relativism

We can also give a Craigean genealogical explanation of the intuitions underlying three standard views of the semantics of knowledge attributions: invariantism, contextualism and (M-)relativism.

A brief explanation of these views will suffice here. There are two ‘dividing lines’ in the literature on the semantics of knowledge attributions. The first concerns the meaning of ‘to know’ (and its cognates). Contextualists hold that uses of ‘to know’ mean different things in different contexts, in much the same way as uses of indexicals like ‘I’. In contrast, invariantists and M-relativists hold that ‘to know’ is univocal: it means the same thing in all contexts. The second dividing line concerns whether knowledge attributions express propositions that are absolutely true or false. While invariantists and contextualists disagree about the contents of the propositions expressed by knowledge attributions they agree that these propositions are absolutely true or false. In contrast, M-relativists hold that the propositions expressed by knowledge attributions are only true or false relative to a context of assessment. Thus, a knowledge attribution may be true as assessed relative to one context, but false as assessed relative to another.

In the literature one can find advocates of all three positions appealing to Craigean genealogy (see Kelp 2011 for invariantism, Henderson 2009 for contextualism and MacFarlane 2014: Ch. 12 for M-relativism). Recall that we are currently interested in genealogical explanations of the *intuitions* underlying philosophical views, not genealogical *vindications*. But the accounts given by Kelp, Henderson and MacFarlane give good indications of how to generate genealogical explanations of their respective positions.

We can start by noting that contextualism fits perfectly with some aspects of Craig’s genealogy. It seems clear that whether a subject is a good informant depends on and varies with the context. If I’m talking to a friend over lunch about Isla’s whereabouts last night and I have good but not conclusive evidence that she was at the party I’ll volunteer myself as an informant on her

whereabouts. However, if I'm giving a statement to the police and I have the same evidence I'll not volunteer myself as an informant on her whereabouts. This suggests that, because what one will require of a good informant depends on and varies with the context, what it means to say someone 'knows' must depend on and vary with the context too. Now, objectivization clearly complicates this picture. But recall Craig's remark that, in some situations, we are still 'in the state of nature'. One way of thinking of objectivization (suggested by Henderson, and by our understanding of Craig) is as pushing us towards co-ordinating our uses of 'to know' in the majority of situations, but allowing us to use the word in ways that are appropriate to our particular situation when the need arises (e.g. when the stakes are particularly high).

However, contextualism fits badly with the central role of testimony in Craig's (and Reynolds' and Williams') genealogy. If uses of 'to know' mean different things in different contexts, then it is not helpful to report that someone else knows (Hawthorne 2004). I may be told that Isla said that Morven knows the bank is open, but unless I also know what epistemic standards Isla was using, this information is useless. Compare: I may be told that somebody said 'I am tired' but, unless I know who said this, this is useless if I want to keep track of who is tired. This suggests that, if we are to have a functioning testimonial practice, we need a univocal semantics for 'to know'.

Enter the M-relativist. MacFarlane thinks that M-relativism can utilise both the aspects of Craig's genealogy that seem to support contextualism and the aspects that seem to support invariantism. The contextualist highlights the fact that we may require different things of good informants in different situations and reasons that these requirements will lead us to use 'to know' in different ways in different situations. But this fact might equally well be taken to support M-relativism: while words like 'to know' (or 'good informant') mean the same thing in all contexts of utterance, whether it is true that someone 'knows' (or is a 'good informant') is relative to the context of assessment. MacFarlane buttresses this point by arguing that relativism is more plausible than contextualism because it places less demands on our memory. He says that, if M-relativism were true, '[t]here would be no need to store a standard with each knowledge attribution, because all of the knowledge attributions would be evaluated in relation to the current standard' (2014: 312). To hammer this point home MacFarlane even offers an 'evolution of assessment sensitivity': it may well have been that 'once upon a time, "knows" behaved just as contextualists say it does' (317). But as social interaction increased and knowledge attributions were exchanged ever more widely across situation and standards, it simply became too tedious to keep track of the standards attached to each attribution. And thus speakers drifted towards the M-relativist understanding.

And yet, while there is a plausible genealogical explanation of why we should expect to have intuitions supporting M-relativism, there is also a plausible story to be told that goes the other way. There are grounds for doubting whether M-relativism is really compatible with the role of testimony in a genealogical account. Consider Reynolds' idea that informants need to internalise the correct testimonial norms. This is difficult to fit with M-relativism: the M-relativist claims testimonial norms require us to assert only propositions that are true relative to our present context of assessment (see MacFarlane 2014: Ch. 12), but if we follow these norms our testimony may be of little use to inquirers in different contexts of assessment. Further, M-relativism poses more general problems for our testimonial practices (Rysiew 2012). We often rely on reports to the effect that someone knows that *p*. But, for the M-relativist, these reports will be true relative to some contexts of assessment and false relative to others. Should the inquirer rely on these reports? Either the M-relativist has to hold that we shouldn't, or they have to allow that we need a way of keeping track of which contexts of assessment reports are true relative to. And thus the seeming advantage of M-relativism over contextualism disappears.

Genealogy and Neutrality

Let us take stock. We have argued that a genealogical account can be given of the intuitions underlying five philosophical views: absolutism, B-relativism, invariantism, contextualism and M-relativism. Some of these views cannot be adopted together. For instance, absolutism and B-relativism are incompatible. This prompts three questions about the relationship between (Craigian) genealogy and relativism:

1. Can genealogy maintain a neutral stance towards these views?
2. Should genealogy maintain a neutral stance?
3. If genealogy cannot (or should not) maintain a neutral stance, which of these should it favour?

We take each question in turn.

Question 1: Genealogy is primarily a *methodology*. Its aim is to make sense of different epistemological 'intuition-pumps' as natural responses to very general social-epistemic situations. On Craig's account, *knowledge* was introduced in order to serve certain deep-seated human needs related to these situations. But it is far from obvious that there should be only one single way in

which these needs can be met. The different genealogical rationales for both absolutism and B-relativism (or for the various views about the semantics of knowledge attributions) make this point vivid: incompatible philosophical positions can be understood as motivated by the same set of general needs. Moreover, in providing genealogical explanations for different, or even opposed, intuitions, we are not necessarily taking a stand on the correctness or rationality of these intuitions. We can understand something as a response to a need without taking the need (or response) to be rational. For instance, you can view a politician's provocative statements as a response to a need for attention without regarding this response as rational, or as a good thing.

A comparison with the "Impartiality" and "Symmetry" principles of the "Strong Programme" of the "Sociology of Scientific Knowledge" (SSK) can help strengthen this point (Barnes and Bloor 1982, Bloor 1991). The Strong Programme 'would be impartial with respect to truth and falsity, rationality or irrationality, success or failure' and it 'would be symmetrical in its style of explanation. The same types of cause would explain say, true and false beliefs.' (Bloor 1991) Thus the proponent of SSK offers the same general types of explanations of theory acceptance in science regardless of whether the respective theories are true or false by our lights. For instance, SSK gives 'symmetrical' sociological explanations of the beliefs of all sides in controversies like the dispute between Boyle and Hobbes over the 'Spring of Air' or the possibility of a mechanically produced vacuum (Shapin and Schaffer 1985).

Genealogy does something similar, albeit at a different level. Where SSK tends to focus on *particular* and contingent historical controversies, the genealogist starts from an *abstract*, idealised and simplified social scenario – the 'epistemic state of nature' -- and then imagines how human needs related to these scenarios lead to the development of one or more concepts. Both forms of analysis are impartial and symmetrical.

And yet, this comparison prompts a worry. Isn't genealogy just an empirically unconstrained form of philosophical speculation? Why engage in genealogy when we have SSK (and related forms of social theory and anthropology)?

At this point we dig in our genealogical heels. The worry is over-stated. First, genealogy is not alone in thinking through abstract and simplified social scenarios. 'Social theory', the abstract and highly theoretical end of social science, does likewise. Take game theory, for example. Few social scientists are inclined to dismiss game theory out of hand just because it abstracts away from details about the social world. Second, genealogy is not entirely unconstrained by empirical information. In Craig's version, we build on information about supposedly universal human needs. In Williams's version, we add further assumptions about human practices, collective goods, free-riding, values,

commitments, and much else. True, Craigean genealogy does not study actual, concrete, dated events. But that is because it investigates ubiquitous and repeatable *types* of events. Finally, recall Williams' insistence that genealogy ultimately has to 'make contact' with actual historical writing. Genealogy is not complete until it has identified the actually endorsed values underlying our epistemic institutions, until it has pinpointed the actual values to which accuracy and sincerity are linked. Going beyond Williams' own writings, one might argue, for instance, that Steven Shapin's study of Boyle's scientific testimony brilliantly brings out the importance of gentlemanly values and conventions surrounding the reporting of facts (Shapin 1994, cf. Kusch 2009).

Question 2: Let us assume that there is a sense in which genealogy can be neutral – it is an explanatory method for making sense of conflicting intuitions. Alas, this does not yet show what exactly would be *valuable* about neutral genealogy. Maybe, if it is to have any interest, a genealogy needs to 'take sides'. The general worry here is that, while an explanation of where our 'intuitions came from' may be of historical or sociological interest, it isn't of philosophical interest.

We are not convinced. If philosophy is anything, it is an exercise of critical self-reflection. And surely, it must be part and parcel of such self-reflection to ask where our (most deeply held) intuitions come from. Their origins may be important for all sorts of reasons. For instance, if Williams is right, understanding the origins of our institution of testimony is crucial for understanding the 'twin values' of sincerity and accuracy. More generally, if we see social epistemology as in some sense continuous with social theory, we surely must care about the social origins of our epistemological intuitions and institutions.

An analogy might be helpful. We see the role of genealogy as being roughly akin to the role of experimental philosophy when it asks which psychological mechanisms and processes influence our judgements. The experimental philosophers do not just claim that it would be interesting to find out what drives our judgements; they claim that their findings are philosophically important. Take, for instance, psychological work on the cognitive processes underlying our judgements about whether subjects 'know' in 'bank cases' (and other similar cases). Some of this work explains intuitions appealed to by contextualists in support of their semantics for knowledge attributions (see e.g. Gerken 2013; Nagel 2008, 2010). The underlying idea is that epistemology is continuous with psychology, and that therefore the epistemologist must care about the psychological origins of our epistemological intuitions. We want to insist that the same goes for epistemology and social theory.

One might object that, while this analogy may clarify the philosophical relevance of genealogy, it does so at the cost of compromising its neutrality. After all, Gerken and Nagel's work is generally

seen as an argument *against* a contextualist account of the semantics of knowledge attributions. We grant that this is both Gerken's and Nagel's intention. But we reject the inference against neutrality. Why should it be assumed that we could not use the experimental method to explain pro-invariantist intuitions too?

Nevertheless, we do not insist on the complete neutrality of genealogy regarding the five positions at issue in this paper. Instead, we shall ask: if genealogy has to take sides, which position should it favour?

Question 3 We are going to focus on the absolutism/B-relativism dispute. The genealogist who 'takes a side' in this dispute seems to have (at least) four options:

- a. The genealogical explanation 'debunks' both absolutism and B-relativism (both are 'empty ideas', cf. Unger 2014).
- b. The explanation leads to philosophical 'quietism' (the question of the 'correctness' of these views never arises).
- c. The explanation 'vindicates' both absolutism and B-relativism.
- d. The explanation vindicates one view, but not the other.

We think the genealogist should plump for d., and this vindicates B-relativism over absolutism.

The problem with a. is that it conflates *origin* and *validity*. It assumes that just because our intuitions have their origins in contingent social settings, our intuitions are worthless. The accusation of worthlessness falls back on this very argument. (This seems like an instance of the 'genetic fallacy' cf. Srinivisan 2015).

The problem with b. is that it is unstable. On one rendering, it collapses into option a.: there is nothing to say because both views have been debunked—they were both wrong, and it was a mistake to have the dispute in the first place. On another reading, it amounts to a form of 'second-order' relativism: there is nothing to be said since both views 'have their place'. But it is hard to see how a resolution of the absolutism/B-relativism dispute that involves second-order relativism could possibly be even-handed. Thus the second reading of b. is really tantamount to d.

Option c. is problematic for the same reason. Absolutism and B-relativism can't 'both have their place': absolutism is committed to non-absolutist views being thoroughly mistaken.

This leaves alternative d. There are two ways of arguing that, if genealogy takes the fourth option, it will end up debunking absolutism. The first is relatively direct. Start with Williams' solution

to the free-rider problem: in order to stop free riders taking advantage of the institution of testimony, we impress upon each other the need to regard accuracy and sincerity as intrinsic, non-negotiable values. Whatever the merits of this solution, it clearly involves a form of anti-realism about the norms of accuracy and sincerity. On Williams' view, these norms are 'of our making', rather than 'being there anyway'. That these norms 'are there anyway' is very much at the heart of the absolutist's credo. Something similar applies to genealogy more generally. Whenever we offer an account of the genealogy of our epistemic concepts and norms, we are presenting them as human constructs, anchored in contingent human needs, subject to the causal influences of biological and cultural evolution. This type of analysis jars with the idea of independent epistemic truths (cf. Street 2009). Admittedly, there is a conceptual difference between anti-realism and B-relativism: while relativism entails anti-realism, the reverse is not true. Nevertheless, anti-realism is incompatible with absolutism. And thus genealogy debunks absolutism. While not being a direct argument *for* B-relativism, it does remove B-relativism's main competitor from the contest.

There is also a second way in which genealogy weakens absolutism. One common line of argument in defence of absolutism—whether in epistemology or in ethics—is that it is the *default* view: the view that we should adopt unless we are given good reason not to (see Cuneo 2007). Part of the reason why absolutism is taken to be the default view is that, if absolutism is true, we can explain a wide range of common intuitions. For example, we can explain why epistemic norms, such as norms of accuracy and sincerity, seem so attractive. Genealogy decisively blocks this route. Genealogical, symmetrical explanations of absolutist *and relativist* intuitions put all of these intuitions on a par. None of these intuitions have the special rights of the firstborn. None of them are the default. Hence the relativist need not fight the uphill battle of challenging allegedly absolutist common sense. Neither side is common sense alone. Both are – up to a point. Given absolutists' reliance on the 'burden of proof', the genealogical redistribution of this burden helps B-relativism.

Finally, where does this leave the invariantism, contextualism and M-relativism dispute? For the reasons just given, it is hard to see how a genealogical account of this dispute could be entirely neutral. But which side does it favour? Answering this question requires getting clear on the relationship between M-relativism and B-relativism: to the extent that M-relativism fits naturally with B-relativism, the argument above offers reasons to think that genealogy favours M-relativism. But this is a complicated issue, which we lack the space to address here.

Summary

In this paper we have brought together two important strands in contemporary epistemology: the debate around epistemic relativism and the genealogy of knowledge. We have proposed – what we believe to be – an original rendering of the latter, and have applied it to the former. We have identified the general human needs and practices that give rise to the intuitions underlying absolutism, B-relativism, invariantism, contextualism and M-relativism. And we have tried to maintain – albeit in a tentative and ‘first-shot’ fashion – that genealogy is a natural ally to B-relativism.²

Bibliography

- Barnes, B. & Bloor, D. (1982), ‘Relativism, Rationalism and the Sociology of Knowledge’, in M. Hollis & S. Lukes (eds.), *Rationality and Relativism*, Oxford: Blackwell, 21-47
- Bloor, D. (1991), *Knowledge and Social Imagery*, 2nd ed., Chicago, Ill.: Chicago University Press.
- Boghossian, P. (2006), *Fear of Knowledge: Against Relativism and Constructivism*, Oxford: Oxford University Press.
- Brown, J. (2006), ‘Contextualism and Warranted Assertibility Manoeuvres’, *Philosophical Studies* 130 (3): 407-435.
- Craig, E. (1990), *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Clarendon Press.
- Craig, E. (1993), *Was Wir Wissen Können*, Frankfurt a. M.: Suhrkamp.
- Craig, E. (2007), ‘Genealogies and the State of Nature’, in A. Thomas (ed.), *Bernard Williams*, Cambridge, UK: Cambridge University Press, 181-200.
- Cuneo, T. (2007), *The Normative Web: An Argument for Moral Realism*, Oxford: Oxford University Press.
- DeRose, K. (2009), *The Case for Contextualism: Knowledge, Skepticism, and Context*, Vol. 1., Oxford: Oxford University Press.
- Gerken, M. (2013), ‘Epistemic Focal Bias’, *Australasian Journal of Philosophy* 91 (1): 41-61.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hazlett, A. (2014). Expressivism and Convention-Relativism about Epistemic Discourse. In A. Fairweather & O. Flanagan (eds.), *Naturalizing Epistemic Virtue*. Cambridge University Press, 223-46.

² Our work on this paper was assisted by funding from the ERC Advanced Grant Project “The Emergence of Relativism” (Grant No. 339382).

- Henderson, D. (2009), 'Motivated Contextualism', *Philosophical Studies* 142 (1): 119-131.
- Kelp, Christoph (2011), 'What's the Point of "Knowledge" Anyway?' *Episteme* 8 (1): 53-66.
- Kusch, M. (2009), 'Testimony and the Value of Knowledge', in A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value*, Oxford: Oxford University Press, 60-94.
- Kusch and McKenna (forthcoming), 'The Genealogical Method in Epistemology'. *Synthese*.
- MacFarlane, J. (2014), *Assessment Sensitivity: Relative Truth and its Applications*, Oxford: Oxford University Press.
- McKenna, R. (2013), "'Knowledge" Ascriptions, Social Roles and Semantics', *Episteme* 10 (4): 335-350.
- Nagel, J. (2008), 'Knowledge Ascriptions and the Psychological Consequences of Changing Stakes', *Australasian Journal of Philosophy* 86 (2):279-294.
- Nagel, J. (2010), 'Knowledge Ascriptions and the Psychological Consequences of Thinking about Error', *Philosophical Quarterly* 60 (2-3): 286-306.
- Reynolds, S. (2017), *Knowledge as Acceptable Testimony*, Cambridge: Cambridge University Press.
- Rorty, R. (1979), *Philosophy and the Mirror of Nature*, Princeton: Princeton University Press.
- Rysiew, P. (2001), 'The Context-sensitivity of Knowledge Attributions', *Noûs* 35 (4): 477–514.
- Rysiew, P. (2012), 'Epistemic Scorekeeping', in J. Brown and M. Gerken (eds.), *Knowledge Ascriptions*, Oxford: Oxford University Press.
- Scheff, T. J. (1988). Shame and conformity: The deference-emotion system. *American Sociological Review*, 53(3), 395-406.
- Shapin, S. and S. Schaffer (1985), *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*, Princeton: Princeton University Press.
- Shapin, S. (1994), *A Social History of Truth: Civility and Science in Seventeenth-Century England*, Chicago: University of Chicago Press.
- Srinivasan, A. (2015), 'The Archimedean Urge', *Philosophical Perspectives* 29 (1): 325-362.
- Street, S. (2009), 'Evolution and the Normativity of Epistemic Reasons', *Canadian Journal of Philosophy* 35: 213-248.
- Unger, P. (2014), *Empty Ideas: A Critique of Analytic Philosophy*, Oxford: Oxford University Press.
- Williams, B. (2002), *Truth and Truthfulness: An Essay in Genealogy*, Princeton and Oxford: Oxford University Press.
- Williamson, T. (2000), *Knowledge and Its Limits*, Oxford: Oxford University Press.