

10 Revealing Social Functions through Pragmatic Genealogies

Matthieu Queloz

Introduction

There is an underappreciated tradition of genealogical explanation that is centrally concerned with social functions. I shall refer to it as the tradition of *pragmatic genealogy*. It runs from David Hume (*T*, 3.2.2) and the early Friedrich Nietzsche (*TL*) through E. J. Craig (1990, 1993) to Bernard Williams (2002) and Miranda Fricker (2007).¹ These pragmatic genealogists start out with a description of an avowedly fictional “state of nature” and end up ascribing social functions to particular building blocks of our practices – such as the fact that we use a certain concept or live by a certain virtue – which we did not necessarily expect to have such a function at all. That the seemingly archaic device of a fictional state-of-nature story should be a helpful way to get at the functions of our actual practices must seem a mystifying proposal, however; I shall therefore endeavor to demystify it in what follows.

My aim in this chapter is twofold. First, by delineating the framework of pragmatic genealogy and contrasting it with superficially similar methods, I argue that pragmatic genealogies are best interpreted as dynamic models whose point is to reveal the function – and noncoincidentally often the *social* function – of certain practices. Second, by buttressing this framework with something it notably lacks, namely an account of the type of functionality it operates with, I argue that both the type of functional commitment and the depth of factual obligation incurred by a pragmatic genealogy depend on what we use the method for: the dynamic models of pragmatic genealogy can be used merely as heuristic devices helping us spot functional patterns, or more ambitiously as arguments grounding our ascriptions of functionality to actual practices, or even more ambitiously as bases for functional explanations of the resilience or the persistence of practices. By bringing these distinctions into view, we gain the ability to distinguish strengths and weaknesses of the method’s application from strengths and weaknesses of the method itself.

The Nature and Point of Pragmatic Genealogies

What are the characteristics of pragmatic genealogies? It quickly becomes apparent that they have little in common with the genealogies approximating regular historiography that one finds predominantly outside philosophy.² Unlike regularly historiographical genealogies, pragmatic genealogies do not in the first instance aim to describe the complex historical roots of a practice and its contingent transformations over the course of history. At the purely formal level which remains largely neutral between different substantive interpretations, the following are typical (though not necessary) features of a pragmatic genealogy:

- i It begins not in a particular time and place, but in a *state of nature*, which differs from the Hobbesian version in that it already contains a small community of social and language-using creatures.
- ii It describes how these creatures have certain *practical needs* which, in the kind of environment they live in, generate *further* practical needs.
- iii In addition, it occasionally also factors in yet further needs that it describes as arising out of *particular sociohistorical configurations* of society.
- iv It shows how this array of needs issues in the need to solve a particular practical *problem*.
- v It presents a particular bundle of dispositions, concepts, institutions, or conventions as a *solution* to this practical problem, thus indicating that given their needs, it would be rational for these creatures to move into a state in which this particular bundle was operative.
- vi It occasionally also describes how the creatures *would be able to get into* that state quite naturally, without much foresight or understanding.

But how are these somewhat quaint-sounding state-of-nature stories best understood? And what are they supposed to tell us? One way of making sense of them is to read them as conjectural depictions of early hominid life in our so-called “environment of evolutionary adaption” (EEA). The state of nature would then be a stand-in for a historical situation about which we have very little data. For just this reason, we would then lack the means to decide whether things *actually* developed as the genealogical story presents them, but we might at least conjecture how things might *possibly* have developed. The state of nature would then be a narrative device by which to give “how possibly” explanations. An example of a genealogical approach that explicitly embraces this interpretation is Philip Kitcher’s *The Ethical Project* (2011).

Where the pragmatic genealogists I listed at the beginning are concerned, however, things must be more complicated. This is because they explicitly deny that their state of nature is in the business of depicting, however conjecturally, the real history of our conceptual practices. Hume writes that the state of nature is “a mere philosophical fiction, which never had, and never cou’d have any reality” (*T*, 3.2.2.14). Williams insists that “the state of nature is not the Pleistocene” (2002, 27), and Craig likewise emphasizes that the “question ‘when?’ just doesn’t apply to it” (2007, 193). Philip Pettit (2018) drives the point home by calling his state of nature “Erewhon,” a Butlerian anagram of “nowhere.”

How, then, are mere philosophical fictions depicting nowhere in particular supposed to tell us anything about reality? The answer comes into view once one sees genealogical fictions as akin to idealizing models in the sciences. Drawing on the interpretation of Craig’s genealogy developed by Martin Kusch and Robin McKenna (Kusch 2009, 2011, 2013; Kusch and McKenna 2018a, 2018b), I suggest that pragmatic genealogies are best interpreted as *dynamic models* – as idealizations with a time axis. These dynamic models start out with a strongly idealized situation, highlight certain practical needs, and show how these needs would drive a community to develop a prototypical form of a certain conceptual practice. The model is then de-idealized towards our actual situation by successively factoring in further needs: needs entailed by the initial needs the model started out with, but also, as in Williams’s case, needs factored into the model based on what we know about the actual history, sociology, and psychology of human beings. But what are these dynamic models for?

The point of these genealogical models is to help us understand what our conceptual practices *do for us*, i.e. whether and how they tie in with our practical needs. To this end, pragmatic genealogies typically work to *reveal social functions* – to render visible unsuspected ways in which our conceptual practices meet social needs. Thus Hume offers a pragmatic genealogy of the virtue of “justice,” i.e. the virtue of respect for property, which exhibits it as a solution to conflict over external goods; Nietzsche presents truthfulness as a solution to the problem of deceit and dissimulation within the community; Craig describes the concept of knowledge as a tool by which to flag good informants and pool information; Williams shows how truthfulness facilitates the gaining and sharing of information; and Fricker highlights the importance of the virtue of testimonial justice in correcting for prejudice. All of these pragmatic genealogists seek to uncover an underappreciated *social function* performed by a conceptual practice – a respect in which it proves beneficial, not (or not just) to the individual who engages in it, but to the social community as a whole.

The genealogical narrative is useful to this end because it perspicuously shows how, from certain practical needs we uncontroversially

have, a need can be derived which we did not necessarily know we had, namely the practical need for the target practice – the practice whose social functionality is to be revealed. In contrast to invisible hand explanations, pragmatic genealogies do not seek to account for the appearance of design by explaining how something obviously functional came to be so; rather, they serve to show how something that does not even *look* designed in fact turns out to be functional in nonobvious ways.³ Using a set of needs which one's addressees identify as having, or which they at least recognize to community to have, we can tell a genealogical narrative showing how, given this set of needs, the target practice is in fact called for by the fact that this set of needs entails a series of further needs issuing in the need for the target practice. This genealogical derivation of needs from needs will then have the form: need A, hence need B, hence need C ... hence need X, where need X is the need for the target practice. Granted that we actually have need A, the genealogical story will then give us reason to think that we also have need X.⁴

It is along just these lines that Williams (2002), for example, seeks to reveal the social function of a community's valuing the truth for its own sake. To value the truth for its own sake, on Williams's view, is in the first instance to value the various states and attitudes expressive of truthfulness, such as accuracy and sincerity, for their own sake (2002, 6–7).⁵ To uncover the point of doing so, he begins with a state-of-nature situation highlighting the need of each individual to gather information about the immediate environment. This is a need that human beings have even on a highly generic conception of the person, which is to say that in ascribing such a need to our agents in the state of nature, we are not expressing a sociohistorically *local* conception of the person (we are not, for instance, expressing a distinctively *liberal* conception of the person by assuming that each individual has a strongly demanding need for autonomy). On the basis of this plausibly generic need-ascription, Williams then points out that the mere fact that individual inquirers occupy different spots at different times already entails that any given inquirer would, under certain circumstances, come to possess what Williams dubs a “purely positional advantage” (2002, 42) over other inquirers. And this means that there are strong pressures on these inquirers not to rely merely on their own senses in acquiring information, but to engage in the practice of sharing or pooling information. But this in turn generates the social need to cultivate in all participating inquirers the dispositions that make good contributors to the information pool: the dispositions which Williams brings under the capitalized headings of “Accuracy” and “Sincerity” to mark the fact that these are, at this point in the story, merely prototypical forms of what we now understand by “accuracy” and “sincerity.” They are as yet merely the dispositions involved in getting one's beliefs right and openly passing them on to others.

Yet as Williams's dynamic model makes clear, development cannot stop there, as the need to cultivate dispositions that make good contributors to the pool itself leads us to consider further needs. The reason is that the dispositions in question cannot deliver what is demanded of them if they do not develop further. The practical value of the individual inquirer's exhibition of these dispositions consists largely in their instrumental value to the community of inquirers: for the individual inquirer, Sincerity is rarely directly of much use, and Accuracy pays only to the extent that its benefits *for the individual* outweigh the costs *to that individual* – but there are many pieces of information that would be of great value to the community even if they are of no particular interest, or involve forbiddingly high risks and dangers, to the individual who could acquire them. Consequently, *mere* dispositions of Accuracy and Sincerity are overly vulnerable to the temptation to free ride, i.e. to profit from the Accuracy and Sincerity of others while failing to exhibit them reliably oneself. And since this holds true for any individual in the community, Accuracy and Sincerity in this form are not stable solutions to the problem of information pooling.

A salient way to overcome this problem, Williams argues, is for Accuracy and Sincerity to come to be regarded as dispositions worth having and worth exhibiting *for their own sake* – they need to come to be regarded (and be commonly known to be regarded) as *intrinsically valuable* dispositions or *virtues*, as Williams puts it (2002, 89–90). And this in turn generates the need for individuals to be capable of *making sense* of these dispositions as intrinsically valuable, which, for Williams, means that their “value must make sense to them from the inside” – they must be able to relate Accuracy and Sincerity “to other things that they value, and to their ethical emotions” (2002, 91–92). If we are to grasp how these prototypical forms of Accuracy and Sincerity have been fleshed out “now and around here,” however, and how they have been “changed, transformed, differently embodied, extended and so on” (Williams 2007, 132) in response to further needs that are more clearly historically and socially situated, the dynamic model must then be de-idealized in that direction. Williams consequently factors in, first, developments in ancient Greece that led to the extension of Accuracy to the distant past (2002, ch. 7); second, developments in eighteenth-century Europe that led to Sincerity's elaboration into the value of authenticity (2002, ch. 8); and lastly, the needs of modern-day liberal democracies to cultivate Accuracy and Sincerity about politics and political history (2002, chs. 9–10, esp. 231–232, 265–266).

Williams's genealogy is not coincidentally about the social function of Accuracy and Sincerity. Such a pragmatic genealogical model serves to sharpen our eye for functional patterns within our actual conceptual practices. If it is to have a point, the functionality it reveals must be one that we are not already fully aware of, and social functions are often

particularly hard to discern. The functionality of practices for the individual are often easier to discern because looking out for individual benefits is something we do anyway, whereas the social point of view is one we only take up in special circumstances or upon reflection. Moreover, Williams's genealogy helps explain *why* the social function of Accuracy and Sincerity is hard to discern: it brings out that the fact that we do not primarily think of Accuracy and Sincerity in functional terms at all – neither in terms of individual nor of social functionality – is an essential part of what renders them functional. Their very functionality demands that their functionality be effaced in favor of intrinsic motivations – which is why, elsewhere, I have called these particular functional dynamics the dynamics of *self-effacing functionality* (Queloz 2018b, §3).

But a pragmatic genealogy derives its interest not just from the fact that it reveals functionality we are not yet aware of; it also derives it from its being something of a mystery how and why the target of the genealogy would have emerged in the first place. Why did individuals ever come to be Accurate and Sincere where it does not pay for them? Yes, this has to do with their valuing Accuracy and Sincerity for their own sake, but to leave it at that is simply to restate the puzzle: why did they ever come to value them intrinsically? Is this more than a fetish, more than a relic from the enchanted world in which God is truth and truth is divine, as Nietzsche (*GM*, III, §24) put it? In answer to these questions, we can usefully construct a model that renders perspicuous how and why Accuracy and Sincerity would arise naturally, without mysterious saltations, in response to individual and social needs – especially if that model can itself explain why these needs would come to include the need for Accuracy and Sincerity to be valued for their own sake, even without metaphysical reasons for doing so. Williams's model to that extent *naturalizes* the intrinsic valuing of Accuracy and Sincerity, and uses this functionally justified intrinsic valuing to explain how a set of practices serving a social function could have arisen *despite* the fact that they were only of limited use to the individuals engaging in them. The genealogy bridges the gap between individual and social functionality on the back of an insight into the social function of intrinsic values.

If pragmatic genealogies are dynamic models revealing social functions, however, it remains unclear which notion of functionality this particular brand of functionalism is supposed to operate with. The pragmatic genealogists themselves provide little guidance in this regard. Their writings leave it underdetermined whether their talk of “functions” should be cashed out in terms of a *causal role* account of functions *à la* Cummins (1975), where functions are ascribed to elements of a system – such as a heating system – on the basis of what they contribute to the realization of some systemic capacity we are interested in (such as the capacity to keep room temperature constant); or in terms of an *etiological* account of functions *à la* Wright (1973), Millikan (1989), and Neander

(2017), where functions are ascribed to items based on what these items were selected to do; or an *agentive functions* account *à la* Searle (2010), where functions are imposed on objects by the purposes of agents, and even biological functions are thought of as causes that serve a purpose relative to agents' values (58–60). On the one hand, it can be seen as a strength of these genealogical approaches that they remain ecumenical regarding the notion of function they involve. This invites one to try and plug in whichever notion one prefers or is interested in. On the other hand, it can also be seen as a weakness, especially in the light of the barrage of objections to functionalist approaches in the social sciences.⁶ These objections express and foster a discomfort with functionalist approaches that is bound to prove an obstacle also to pragmatic genealogies once their functionalist spirit is recognized. In the remainder of this essay, I will therefore try to dispel such discomfort by laying bare the functionalist innards of pragmatic genealogies.

Functions in and Functions of Pragmatic Genealogy

The key to understanding how pragmatic genealogies work is to distinguish what they minimally and primarily do from what they can then be used to do on that basis. In the first instance, pragmatic genealogies serve to reveal instrumental relations between certain needs and certain conceptual practices *within a fictional model*. They issue in conclusions of the form: “The function of the prototype of X is to satisfy a need to Y.” For the pragmatic genealogists discussed above, this formula yields the following claims:

- The function of the prototype of the virtue of justice is to satisfy a need to avoid conflicts over external goods (Hume).
- The function of the prototype of the virtue of truthfulness is to satisfy a need to avoid deception within the community (Nietzsche).
- The function of the prototype of the concept of knowledge is to satisfy a need to flag good informants (Craig).
- The function of the prototype of the virtues of Accuracy and Sincerity is to satisfy a need to gain and share information effectively (Williams).
- The function of the prototype of the virtue of testimonial justice is to satisfy a need to correct for prejudice (Fricker).

Each of these claims involves an *ascription of functionality*: it highlights an instrumental relation between a prototypical concept or virtue on the one hand and a need as represented in the model on the other. This does not yet in itself carry any claims about our actual situation, past or present. But the point of working with such a model is to *render visible* similar instrumental need-concept or need-virtue relations in the actual history of our conceptual practices or within our current conceptual

practices. The dynamic model that is the pragmatic genealogy constitutes a useful guide to the discernment of such functional patterns because it represents them in their clearest, most generic form, free of the clutter and complexities of reality, thus sharpening our eye for similar functional patterns in reality – much as a priming look at a prototypical morel will assist the morel seeker in spotting morels of varying shapes and colors hidden underneath the tangle of twigs. And to the extent that we indeed find such instrumental relations in the past or the present, we can use that as a basis for functional explanations or functional assessments: insofar as the need-concept or need-virtue relation obtained in the past, this helps explain why we came to have the concept or virtue; insofar as the need-concept or need-virtue relation obtains now and around here, this gives us reasons for or against continuing to cultivate the concept or virtue.

Moreover, the dynamic models provide *prima facie* evidence for these functionality ascriptions in much the same way that design analyses in evolutionary biology do: in both cases, a model is used to show that a given trait would solve a problem, and this is used as evidence for thinking that the traits we actually find actually solve similar problems (see Kincaid (1996, 118–119) for a discussion of design analyses in biology). An important difference is that in biology, design analyses typically turn on *optimality arguments* – arguments to the effect that a given trait, such as a certain foraging strategy or a reproductive strategy, can be mathematically shown to form an optimal solution to a problem. Pragmatic genealogies, by contrast, turn on what might be called *indispensability arguments*. These can be very roughly characterized as taking the following form:

- P1 Given certain root needs RN , creatures in the state of nature would need to solve a certain problem.
- P2 A conceptual practice P with characteristics C would form a salient solution to that problem.
- P3 There is an undemanding path by which creatures in the state of nature could come to develop conceptual practice P .
- P4 We share something like the root needs RN , namely RN^* , and we also have something like practice P , namely P^* .
- C1 Therefore, P^* likely functions to satisfy RN^* .

This characterization still begs numerous questions, including notably questions about how much like us these creatures and their environment are, and what the criteria for similarity between RN and RN^* and P and P^* should be – I say more on these questions in Queloz (Forthcoming-b). Moreover, it should be noted that (P3) does not figure prominently in all pragmatic genealogies, although it is an important and much celebrated feature of Hume's genealogy that he traces out a path by which the virtue of justice might arise without much foresight or explicit coordination.⁷ But the point I want to press here is that this line of argument does not

involve the idea that some behavioral trait can be mathematically shown to be optimal. The line of argument turns on the idea that a behavioral trait – more specifically, a concept or virtue – of *some* broadly outlined form, characterized only by its ability to discharge the function at issue, constitutes an indispensable solution to a problem that any creature with certain needs faces. The argument does not purport to show that *P* is optimal; it purports to show that *P* is conditionally *necessary*.

On this account, the core of a pragmatic genealogy is an ascription of functionality relative to needs. Drawing on a certain understanding of what human beings need, certain concepts or virtues are presented as having, among their various actual and potential effects, such effects as tend either directly or indirectly to meet those needs. As Kincaid (1990, 1996, 2006, 2007) and others have argued, ascriptions of this sort are innocuous and capable of surviving scrutiny by those who are skeptical of functionalism in the social realm, because such ascriptions of functionality are straightforwardly identifiable with a set of causal claims: the practice of using or living by a concept or a virtue has certain effects; these effects contribute to the satisfaction of certain needs; and, given the transitivity of causation, the practice of using or living by the concept or virtue therefore helps satisfy these needs.

With any ascription of functionality, there is a question about the extent to which the functionality in question is *observer-relative*; to what extent is this the case for the functionality ascriptions of pragmatic genealogy? Are they more like Searle's agentic functionality ascriptions, which are dependent on the purposes agents happen to have? Or are they more like etiological ascriptions of "proper functions," in Millikan's terminology, which achieve independence from subjective purposes by basing ascriptions of functionality on objective selection histories?

There are reasons to think that pragmatic genealogies are interestingly situated between these two poles. On the one hand, they involve an observer-relative dimension insofar as they take their basic normative orientation, which all functional talk requires in one way or another if it is to allow some sort of discrimination between the functional and the dysfunctional,⁸ from *needs*, and need ascriptions involve a substantial degree of interpretation: needs are something that an entity lacks as long as it is described merely in the vocabulary of physics; it is only once it is described in more normative terms that needs come into view. At the minimum, these terms must permit a distinction between survival and death, and more richly normative terms might allow for further distinctions, first between bare survival and flourishing, and then between flourishing by, in Williams's phrase, "the ethological standard of the bright eye and the gleaming coat" (2011, 52) and more demanding (but also more socio-historically local) standards of flourishing, which may include such things as the need for autonomy or political self-determination. On the other hand, needs are also more objective than

ends and purposes. Unlike ends or purposes, needs are something one cannot *decide* to have – they are not subject to one’s will. Nor are needs subject to the will of the observer, just as their presence or absence does not depend on the observer’s ends or purposes. Furthermore, needs are not luminous, which is to say that one can have them without knowing that one has them. This is also part of what gives pragmatic genealogies their point, because they can reveal that we have certain needs which we may not have been aware we had. Hence, whether or not we have certain needs – such as a need for the virtues of Accuracy and Sincerity, or a need for the concept of knowledge – is, in all these respects, an objective matter and something that can come as a discovery.

Functionality ascription, then, is the primary business of pragmatic genealogy; but once a pragmatic genealogy has suggested and buttressed a functionality ascription, this ascription can be used as an explanatory basis for functional explanations of varying ambitions. One thing that this explanatory basis might be used for is to explain the stability or, in Philip Pettit’s (1996, 2000) terminology, the *resilience* of the building blocks of our way of life:

Which are the more or less passing ephemera and which the phenomena that are deeply embedded in the society? Which are more or less incidental or contingent features and which are features apt to last? There is an interesting research programme suggested by such questions. It would take any society or culture or institution and, reviewing the data on various traits displayed by the entity in question, would seek to separate out the dross from the gold. It would try to identify and put aside the features that may be expected to come and go. And it would seek to catalogue the more or less necessary features that the society or culture or institution displays. It would give us a usefully predictive stance on the society, providing us with grounds for thinking that such and such features are likely to stay, such and such other features likely to disappear.

(Pettit 1996, 299–300)

The research program described here might well be pursued using pragmatic genealogy as one’s method. Note, however, that this research program need not be backwards-looking at all: claims of resilience need not involve any factual claims about history being a certain way. To say that a concept or virtue is resilient because it stands in an instrumental relation to needs is not necessarily to say that it came to exist for that reason, or even that this instrumental relation played any role in its historical development. The conceptual practice could just have popped into existence, or it might have been instituted by a mad king on a whim, and yet it might truly be said to be resilient for functional reasons, i.e. in ways that could be explained by reference to the instrumental relations

to needs that the conceptual practice, whatever its genesis, now stands in. Pettit (2000, 48) gives the example of golf clubs, asking us to imagine that they have come into existence purely because people enjoy the motions involved in playing golf. “Consistently with the absence of any such historical selection,” he notes, “what might well be the case is that golf clubs have certain effects, certain functional effects, such that were they to come under any of a variety of pressures, then the fact of having those effects would ensure that they survived the pressure,” and if so, we can conclude that “though not the beneficiaries of actual selection, golf clubs do enjoy the favor of a virtual process of selection” (2000, 48).

Yet to say that the effects of golf clubs would ensure that they survived any of a variety of pressures is no doubt too strong, since it would imply that golf clubs are not just resilient, but, as one might put it, *hyper-resilient*. The more modest claim advanced by a resilience explanation should rather be that *if*, as a community, we were to move away from golf clubs, this would rob us of some of their functional effects, and this loss would make itself felt (at least among those for whom golf clubs are presumed to have functional effects; the fact that this is likely to be a rather exclusive group, and that the functions in question are perhaps not functions anyone outside that group would want to see discharged in the first place, indicate further reasons why the ambitions of resilience explanations are better reined in). The conclusion we then reach is that there would be some pressure, however limited in strength and scope, to maintain golf clubs – though whether there would be enough pressure for them ultimately to *survive* in the face of countervailing pressures remains a further question.

A second type of explanation that a functionality ascription might be used for is the explanation of the actual historical *persistence* of a concept or virtue. This still stops short of claiming that it *came into existence* because it serves a need (a claim that would raise questions about how as yet unrealized effects can bring something to exist). But it does involve committing oneself to the claim that *the fact that the concept or virtue endured* or was retained once it had come into existence had something to do with its relation to our needs.

In sum, a pragmatic genealogy is in the first instance a narrative device by which to reveal and ground ascriptions of functionality which can then be used to explain the resilience or even the persistence of concepts or virtues. We thus get the following schema, with (1)–(3) together potentially acting as the explanatory basis for an explanation either of the form of (4) or of the form of (5):

Functionality Ascription (Explanatory Basis)

- 1 The practice of living by concept/virtue *A* causes (sometimes via inferential consequences) the consequence *B*.

- 2 *B* helps satisfy need *C*.
- 3 The practice of living by *A* helps satisfy *C*.

Functional Explanation

- 4 *A* is resilient because it helps satisfy *C*.
- 5 *A* persists because it helps satisfy *C*.

The key idea that this schema is meant to bring out is that pragmatic genealogy minimally and primarily serves to substantiate (1)–(3), and it *can*, though it *need not*, also serve to substantiate (4) or (5), be it with regard to our present situation, with regard to some historical situation, or both. Two aspects of the schema bear further clarification, however. First, it is true that on a possible reading of (5), (5) claims no more than (4): in Kincaid’s (2007, 223) rendering, for example, one way for something to persist because it has certain consequences is for it to be *resilient now*, whatever its history, in virtue of the fact that there *would* be pressures to keep it *if* we moved away from it. By contrast, I distinguish (4) and (5) precisely to register the fact that we do undertake commitments about the history of the concept or virtue as soon as (5) is read as implying not just *counterfactual* claims, i.e. claims to the effect that certain forces *would* be actualized *if* a conceptual practice were to diverge from a certain functionally specified configuration, but *factual* claims to the effect that such forces *were* actualized and are part and parcel of the causal-historical story explaining why we now find the concept or virtue.

Second, the qualification “sometimes via inferential consequences” in (1) registers the fact that the relevant practical consequences of living by a concept can be more or less immediate consequences of particular acts of concept application. Sometimes, the relevant practical consequence of living by a concept is simply the fact that this renders the concept-user suitably sensitive to the presence of certain items in the world; but sometimes, the relevant consequences lie further downstream, and the path towards them might lead through inferential consequences, i.e. through the inferences the concept-user is put in a position to draw by coming to live by the concept. Michael Dummett (1973, 454) gives the example of a student learning the concept of validity, and thus coming to be able to distinguish between valid and invalid arguments. But what gives the concept of validity its *point* is not the capacity to make this distinction in itself, but rather the practical difference made by the *inference one can then draw* from an argument’s being valid, namely that one has reason to accept the conclusion given that one accepts the premises. Consequently, a student who reliably applied the concept but never drew this inference – who treated the distinction between valid and invalid arguments as being like the distinction between Petrarchan and

Shakespearean sonnets – could perhaps be said to have acquired the concept of validity, but when used in this way, the concept would fail to tie in with anyone’s needs.

We can illustrate the schema with Craig’s genealogy. Craig’s dynamic model highlights the way in which social and language-using creatures would be driven to develop something like our concept of knowledge by two sets of practical pressures: the first set of practical pressures grows out of the fact that each individual needs information about his or her immediate environment, and is to that extent in the position of the *inquirer*: someone who wants to find out *whether p*. But for social and language-using beings like us, there are strong incentives to rely not just on one’s own senses in acquiring information, but to tap into others’ stores of information. This means that there are pressures on each inquirer to become able to identify what, given the particular needs and capacities of that inquirer, are good informants as to whether *p*. This the inquirer becomes able to do by developing the concept of what Kusch (2009, 65) aptly calls *proto-knowledge* – a concept that serves to *flag good informants*.

Proto-knowledge is still markedly different from our concept of knowledge, however, in that it remains strongly indexed to the situation of the individual concept-user. It tracks whomever is a good informant *for me*, given *my* needs and capacities, *here* and *now*. But if we factor in the second set of practical pressures, we come to see why we in fact operate the concept of knowledge rather than the concept of proto-knowledge. This second set of practical pressures arises from the fact that inquirers have a strong interest in *recommending informants to each other*; and the more they do – the more they socially cooperate not just in exchanging information, but in exchanging information about who is a good source of information on a given question – the more they have reason to operate a concept that is less subjectivized than *proto-knowledge*. It is this second set of pressures which leads to the concept of someone who is a good informant whether *p* for anyone, whatever their needs and capacities, anywhere and at any time: someone, in other words, who *knows* whether *p*. The concept of knowledge we end up with in Craig’s model is thus revealed to perform a social function that is central to a kind of epistemic division of labor, namely the social practice of information pooling.

Plugging Craig’s genealogy of the concept of knowledge into the schema then yields the following:

Functionality Ascription (Explanatory Basis)

- 1 The practice of living by the concept of knowledge causes the flagging of good informants.
- 2 The flagging of good informants helps satisfy the need to pool information.

- 3 The practice of living by the concept of knowledge helps satisfy the need to pool information.

Functional Explanation

- 4 Were the practice of living by the concept of knowledge to come under pressure, there would be some pressure to drive it back into use because it helps satisfy the need to pool information.
- 5 The fact that the practice of living by the concept of knowledge helped satisfy the need to pool information in the past caused the concept to be retained.

To which of these claims does Craig's genealogy commit him? As he himself points out, this depends on the purposes which the genealogical story is taken to serve. "The depth of factual obligation incurred by a state-of-nature theory depends on its aims," he writes; it "will be greatest when its intentions are explanatory, to account for the existence of the target phenomenon" (Craig 2007, 193). This underscores the important methodological point that pragmatic genealogies *in themselves* – i.e. the bare genealogical narratives considered in isolation from the context and spirit in which they are advanced – do not yet determine how much is being claimed and what evidence they are beholden to. These parameters only receive determinate values once the pragmatic genealogy is put to use in a particular context with a view to performing a particular task, and these will be different values in different contexts.

In Craig's own case, the aim is in the first instance to cure us of the temptation to define our present concept of knowledge in terms of necessary and sufficient conditions by getting us to look at the concept from a pragmatic point of view. Once we let our understanding of the concept grow genealogically out of our understanding of the needs of inquirers, we will understand why we should not expect the application conditions of the concept of knowledge to be necessary conditions at all, but rather to reflect the typical conditions under which the concept has a point for inquirers given their needs. As we saw, this need in principle amounts to no more than to offering a model of a need-concept relation as a heuristic device by which to reveal whether such relations also obtain in past or present societies. But Craig aims to do more than that. In his later reflections on *Knowledge and the State of Nature*, he emphasizes that he "was trying to *explain how certain real results have arisen*, and only real pressures can produce real results" (Craig 2007, 190, emphasis mine). In accordance with this aim, he notes:

I do and must suppose that there were societies whose members, collectively and individually, had the needs I ascribe to them and were able, whether as the outcome of some conscious process or of other equally real tendencies, to find their way to the solution I describe. ...

My line was, and had to be, that the needs were real and the persons concerned would have come, in one way or another, to satisfy them. ... I had to maintain that the circumstances that favour the formation of the concept of knowledge still exist, or did until very recently, since otherwise I would have had no convincing answer to the obvious question why it should have remained in use

(Craig 2007, 191)

On this Craigean reading of Craig, the pragmatic genealogy does not just serve as a device by which to sharpen our eye for (1)–(3); nor does it rest content to claim that (1)–(3) are now, for whatever reason, the case, and that therefore the concept of knowledge is now resilient in the sense of (4); rather, the genealogy is used to claim (5): to explain why the concept of knowledge persists.

Of course one might invoke the Death-of-the-Author principle and insist that authorial intentions are not always the most reliable guide to a book, but if the account of pragmatic genealogy offered here is along the right lines, we might accept Craig's self-interpretation, find it insufficiently corroborated by evidence and short on detail and mechanisms, and yet think no worse of his pragmatic genealogy, because the merits of Craig's genealogy are distinct and separable from the use he made or took himself to make of it. We must distinguish the method – offering a dynamic model of the functional relations between needs and conceptual practices – from the use to which it is put. If, like the Craig of 2007, we harbor the ambition to use the genealogy outlined in *Knowledge and the State of Nature* to explain why the concept of knowledge *persists*, we will have incurred a different type of functional commitment – namely an etiological one – and correspondingly deeper factual obligations than if we used it simply to reveal the relation of the concept we now have to some of our present needs. Using the genealogy effectively to explain the persistence of the concept of knowledge might well require rather more supporting material than Craig had room for in his dense and admirably concise book. But the important point is that this does not invalidate the genealogy. The genealogy itself is merely a multipurpose model, a tool that earns its keep in many trades.

Our conclusion, then, is that both the *type* of functional commitment and the *depth* of factual obligation incurred by a pragmatic genealogy *depend* on what we use it *for*. We can use the dynamic models of pragmatic genealogy merely as heuristic devices by which to sharpen our understanding of how certain needs bring certain problems and call for certain solutions while retaining an open mind as to whether a given society exhibits either those needs or anything like the modelled answer to them; the model would then serve to sharpen our eye for the needs and the functional patterns they tend to engender. But we can also use pragmatic genealogy to reveal (and to bolster our case for)

factual claims about the present or the past. On this basis, we might make backward-looking use of pragmatic genealogy, deploying it to account for the persistence and ubiquity of certain arrangements on the basis of what the genealogy reveals to be their social function. But we might equally make forward-looking use of pragmatic genealogy, either confidently to predict that certain arrangements are not going to go away, because the genealogy reveals them to be resilient; or, on the contrary, to highlight the reasons we have to cherish and defend these arrangements against countervailing pressures because the genealogy reveals unsuspected ways in which the satisfaction of certain needs depends on them. Pragmatic genealogy can thus not only be used to show what something *would do for us* if we were in a fictional state of nature; but also what it has done, what it now does, and what it can continue to do.

Notes

- 1 More recent additions to the tradition include Martin Kusch and Robin McKenna (2018b) and Philip Pettit (2018).
- 2 Pragmatic genealogies also differ from counterfactual state-of-nature scenarios in a Hobbesian vein, because in the latter, the state of nature is usually interpreted as playing a primarily *justificatory* role in an argument to the effect that a situation in which there was some form of public authority would be superior even to the best anarchical state of nature one could reasonably hope for (see Nozick 1974, 5; Kavka 1986; Hampton 1987). Nozick's own use of state-of-nature stories is more complicated, as he does harbor explanatory ambitions. But see Williams (2002, ch. 2) for a discussion of why these raise problems of their own.
- 3 For a discussion of invisible hand explanations, see Ullmann-Margalit (1978, 1997) and Tieffenbach (2011, 2013). For a comparison with genealogical explanations, see Williams (2002, 31–32, 253).
- 4 I say more about when and why one would want to resort to dynamic models to reveal functionality in Queloz (2017, 2019a, 2019b, Forthcoming-a).
- 5 For a more detailed discussion of Williams's genealogy, see Queloz (2018).
- 6 I am thinking here of charges such as the following: that functionalism draws on mysterious forms of teleology by ascribing free-floating purposes without tying these back to agents or minds whose purposes they are; that it uses future effects to explain present causes; that its explanations are missing an underlying mechanism; that ascriptions of functionality are unfalsifiable; that they are vacuous; that they draw illegitimate analogies to biology; that they fail to tell us why we have the practices we have rather than equally useful alternatives. For rebuttals of all these charges using accounts of functions along the lines of those I draw on here, see in particular Kincaid (1996, 2007), Pettit (1996, 2000), Barnes (1995, 2000), and Rosenberg (2016a, 2016b, ch. 10). Historically, the work of Ayala (1970), Wimsatt (1972), Wright (1973), Cummins (1975), Cohen (1978), Brandon (1981), Millikan (1984), and Neander (1991) proved seminal in rehabilitating appeals to functionality. For helpful overviews of how the notion of a function is understood in the philosophy of biology, see Buller (1999) and Ariew, Cummins, and Perlman (2002).

- 7 For discussions of Hume's genealogy, see Baier (1988), Blackburn (2008), and Cohon (2008). For discussions of its subsequent importance to game theory, see Binmore (2005, 4), Charron (1980), and Lewis (2002, 4).
- 8 Functions, as Millikan aptly puts it, are "a measure from which actual facts can depart" (2005, 83).

References

- Ariew, Andre, Robert Cummins, and Mark Perlman. 2002. *Functions: New Essays in the Philosophy of Psychology and Biology*. Oxford: Oxford University Press.
- Ayala, Francisco J. 1970. "Teleological Explanations in Evolutionary Biology." *Philosophy of Science* 37, no. 1: 1–15.
- Baier, Annette. 1988. "Hume's Account of Social Artifice: Its Origins and Originality." *Ethics* 98, no. 4: 757–78.
- Barnes, Barry. 1995. *The Elements of Social Theory*. Cambridge: Cambridge University Press.
- . 2000. *Understanding Agency: Social Theory and Responsible Action*. London: SAGE.
- Binmore, Ken. 2005. *Natural Justice*. New York: Oxford University Press.
- Blackburn, Simon. 2008. *How to Read Hume*. London: Granta.
- Brandon, Robert N. 1981. "Biological Teleology: Questions and Explanations." *Studies in History and Philosophy of Science* 12, no. 2: 91–105.
- Buller, David J. 1999. *Function, Selection, and Design*. New York: SUNY Press.
- Charron, William C. 1980. "Convention, Games of Strategy, and Hume's Philosophy of Law and Government." *American Philosophical Quarterly* 17, no. 4: 327–334.
- Cohen, G. A. 1978. *Karl Marx's Theory of History: A Defence*. Princeton, NJ: Princeton University Press.
- Cohon, Rachel. 2008. *Hume's Morality: Feeling and Fabrication*. New York: Oxford University Press.
- Craig, Edward. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- . 1993. *Was wir wissen können: Pragmatische Untersuchungen zum Wissensbegriff*. Wittgenstein-Vorlesungen der Universität Bayreuth, edited by Wilhelm Vossenkuhl. Frankfurt am Main: Suhrkamp.
- . 2007. "Genealogies and the State of Nature." In *Bernard Williams*, edited by Alan Thomas, 181–200. Cambridge: Cambridge University Press.
- Cummins, Robert C. 1975. "Functional Analysis." *Journal of Philosophy* 72: 741–764.
- Dummett, Michael. 1973. *Frege: Philosophy of Language*. New York: Harper and Row.
- Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford; New York: Oxford University Press.
- Hampton, Jean. 1987. *Hobbes and the Social Contract Tradition*. Cambridge: Cambridge University Press.
- T: Hume, David. 2000. *A Treatise of Human Nature*, edited by David Fate Norton and Mary J. Norton. Oxford: Oxford University Press.
- Kavka, Gregory S. 1986. *Hobbesian Moral and Political Theory*. Princeton, NJ: Princeton University Press.

- Kincaid, Harold. 1990. "Assessing Functional Explanations in the Social Sciences." *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1990: 341–354.
- . 1996. *Philosophical Foundations of the Social Sciences: Analyzing Controversies in Social Research*. Cambridge: Cambridge University Press.
- . 2006. "Evolutionary Social Science Beyond Culture." *Behavioral and Brain Sciences* 29, no. 4: 356–356.
- . 2007. "Functional Explanation and Evolutionary Social Science." In *Philosophy of Anthropology and Sociology*, edited by Stephen P. Turner and Mark W. Risjord, 213–247. Amsterdam: North-Holland.
- Kitcher, Philip. 2011. *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Kusch, Martin. 2009. "Testimony and the Value of Knowledge." In *Epistemic Value*, edited by Adrian Haddock, Alan Millar and Duncan Pritchard, 60–94. Oxford: Oxford University Press.
- . 2011. "Knowledge and Certainties in the Epistemic State of Nature." *Episteme* 8, no. 1: 6–23.
- . 2013. "Naturalized Epistemology and the Genealogy of Knowledge." In *Contemporary Perspectives on Early Modern Philosophy: Nature and Norms in Thought*, edited by Martin Lenz and Anik Waldow, 87–100. Dordrecht; New York: Springer.
- Kusch, Martin, and Robin McKenna. 2018a. "The Genealogical Method in Epistemology." *Synthese*. doi:10.1007/s11229-018-1675-1.
- . 2018b. "The Genealogy of Relativism and Absolutism." In *Metaepistemology: Realism and Anti-Realism*, edited by Christos Kyriacou and Robin McKenna, 217–239. New York: Palgrave Macmillan.
- Lewis, David. 2002. *Convention: A Philosophical Study*. Oxford: Blackwell.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA; London: MIT Press.
- . 1989. "In Defense of Proper Functions." *Philosophy of Science* 56, no. 2: 288–302.
- . 2005. *Language: A Biological Model*. Oxford: Clarendon Press.
- Neander, Karen. 1991. "Functions as Selected Effects: The Conceptual Analyst's Defense." *Philosophy of Science* 58, no. 2: 168–184.
- . 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. Cambridge, MA: MIT Press.
- TL: Nietzsche, Friedrich. 1979. "On Truth and Lies in a Nonmoral Sense." In *Philosophy and Truth: Selections from Nietzsche's Notebooks of the Early 1870's*, edited and translated by Daniel Breazeale, 79–97. Atlantic Highlands, NJ; London: Humanities Press International.
- GM: ———. 1998. *On the Genealogy of Morality*, translated by Maudemarie Clark and Alan J. Swensen. Indianapolis, IN; Cambridge: Hackett.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. Oxford: Blackwell.
- Pettit, Philip. 1996. "Functional Explanation and Virtual Selection." *British Journal for the Philosophy of Science* 47, no. 2: 291–302.
- . 2000. "Rational Choice, Functional Selection and Empty Black Boxes." *Journal of Economic Methodology* 7, no. 1: 33–57.
- . 2018. *The Birth of Ethics: Reconstructing the Role and Nature of Morality*, edited by Kinch Hoekstra. Oxford; New York: Oxford University Press.

- Queloz, Matthieu. 2017. "Nietzsche's Pragmatic Genealogy of Justice." *British Journal for the History of Philosophy* 25, no. 4: 727–749.
- . 2018. "Williams's Pragmatic Genealogy and Self-Effacing Functionality." *Philosophers' Imprint* 18, no. 17: 1–20.
- . 2019a. "Genealogy and Knowledge-First Epistemology: A Mismatch?" *The Philosophical Quarterly* 69. doi:10.1093/pq/pqy041.
- . 2019b. "The Points of Concepts: Their Types, Tensions, and Connections." *Canadian Journal of Philosophy* 49, no. 8: 1122–1145.
- . Forthcoming-a. "From Paradigm-Based Explanation to Pragmatic Genealogy." *Mind*. doi:10.1093/mind/fzy083.
- . Forthcoming-b. "How Genealogies Can Affect the Space of Reasons." *Synthese* Online First. doi:10.1007/s11229-018-1777-9.
- Rosenberg, Alex. 2016a. "Functionalism." In *The Routledge Companion to Philosophy of Social Science*, edited by Lee McIntyre and Alex Rosenberg. London: Routledge. 147–157.
- . 2016b. *Philosophy of Social Science*. Boulder, CO: Westview Press.
- Searle, John. 2010. *Making the Social World: The Structure of Human Civilization*. Oxford; New York: Oxford University Press.
- Tieffenbach, Emma. 2011. "Invisible-Hand Explanations." PhD diss., University of Geneva.
- . 2013. "Invisible-hand Explanations: From Blindness to Lack of Witness." *Social Science Information* 52, no. 3: 450–470.
- Ullmann-Margalit, Edna. 1978. "Invisible-Hand Explanations." *Synthese* 39: 263–291.
- . 1997. "The Invisible Hand and the Cunning of Reason." *Social Research* 64, no. 2: 181–98.
- Williams, Bernard. 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton, NJ: Princeton University Press.
- . 2007. "Truth and Truthfulness." In *What More Philosophers Think*, edited by Julian Baggini and Jeremy Stangroom, 130–146. London: Continuum.
- . 2011. *Ethics and the Limits of Philosophy*. Routledge Classics. London; New York: Routledge.
- Wimsatt, William C. 1972. "Teleology and the Logical Structure of Function Statements." *Studies in History and Philosophy of Science Part A* 3, no. 1: 1–80.
- Wright, Larry. 1973. "Functions." *Philosophical Review* 82, no. 2: 139–168.