The Function of Morality

1 INTRODUCTION

What is the function of morality? On this question, something approaching a consensus has recently emerged. Many philosophers now tell us that the function of morality is to reduce social tensions, and to thereby enable a society to smoothly and efficiently ensure the well-being of its members.

I think that it is time to subject this consensus to more rigorous scrutiny. My view is that the collection of practices, beliefs and dispositions we call 'morality' is far more functionally complex than the standard story would have us believe. Morality may indeed reduce social tensions in certain contexts, but it may also inflame them in others, and it probably plays a variety of other distinct roles in human societies. None of these, I think, represents the primary function of morality. However, I cannot argue for this positive view here. Instead, my task in this paper is primarily negative: I will attack the evidentiary relation between an evolutionary *genealogy of morals* and the functional claim itself. I want to ask: suppose morality *did* emerge as the standard genealogy says that it did; could that story provide the basis for beliefs about morality's current function? As I hope to show, it is far from clear that it can.

My argument is thus a contribution to the wider debate over the role of evolutionary biology in explanations of human social life. While some theorists have wished to draw on ancient evolutionary history in order to motivate substantive conclusions about the function(s) of human social practices, I will show that this type of inference runs into a serious and possibly insurmountable epistemological problem.

I will begin by describing the functional claim itself before outlining the genealogical inference that often underlies it. I will then pose a problem for those who might wish to derive the function of morality from its ancient history, and conclude with a few thoughts on how functional claims of this sort might be established in a more responsible fashion.

2 THE FUNCTION OF MORALITY

The following is a representative set of quotations from those who cleave to the functional hypothesis in question:

> Why do we have moral values? The obvious answer is that morality emerges
> as a system of rules for getting people to function collectively in stable and

productive ways. We have morality to build a coherent social group. Moral values lead us to cooperate and prevent us from harming members of our communities (Prinz 2007, p.32).

Moral thinking has a function, I have argued—both evolutionarily and contemporarily…. The *moralization* of our practical lives contributes to the satisfaction of our long-term interests and makes for more effective collective negotiation by supplying license for punishment, justification for likes and dislikes, and bonding individuals in a shared framework of decision-making (Joyce 2006, p.208).

[T]he evolutionary function of moral judgments… is *interpersonal coordination*. Roughly, moral judgments are the products of a mechanism that allows groups of interacting individuals to co-ordinate their actions and emotions for mutual benefit. The function of the moral habit is therefore to produce mutually beneficial co-operative patterns of action and emotion (Sinclair 2012, p.14).

The function of morality is the enhancement of social cohesion via the amplification of our psychological altruistic dispositions (Kitcher 2005, p.178).

The gloss varies, but the underlying idea is very similar. For each of these philosophers, morality functions to reduce interpersonal conflict in a way that enables people to mutually flourish.

Now, one easy way to attack such claims is to deny that there is a coherent, unified, historically stable phenomenon for us to study at all. I am sympathetic to those who worry that the category 'morality' is coarse-grained in the extreme, and also with those who complain that the reference of the term seems to change depending on which literature one engages with.[1] That said, in this paper, I will assume that this problem can be solved, and that there is some reason to treat 'morality' as a relatively unified, coherent, and historically stable phenomenon. Of course, a complete theory of this sort will require careful distinctions between moral behavior, moral language and moral judgment, especially since each of these phenomena might have distinct functional profiles.[2] However, since these phenomena are very closely related, I will continue to use the term 'morality' to refer

---

[1] For example, the evolutionary-psychological literature displays a strong tendency to simply equate morality with some form of altruism. See (De Waal 2013). Philip Kitcher is more cautious, including certain capacities for normative guidance along with characteristic emotional responses (Kitcher 2011, ch.2).

[2] I thank an anonymous reviewer for this important observation.

to a certain cluster of dispositions, feelings and judgments, and I will assume that we have an intuitive grasp of the phenomenon under study.[3]

Now, it might be that some of the philosophers quoted above would, if pressed, distinguish the ancient or evolutionary function of morality from its *current* function, that is, from the dispositions and characteristic effects it has in present-day societies. Nothing in this paper is meant to undermine claims about morality's ancient functions. Rather, I am interested in a specific inference, namely, one which moves from the evolutionary function to the contemporary one.

After all, the present-day functional claim at issue here is often, though not always, derived from an etiology of moral belief and practice. Kitcher is most explicit about the assumption which underlies the inference:

> [I conceive] the current human situation as analogous to that initially prompting the ethical project. As it was in the beginning, so too now—for the conflicts to which our ancestors' lives were subject are mirrored in contemporary hostilities across the human population. According to this vision, the original function of ethics—to remedy altruism failures— remains primary. (Kitcher 2011, p.8)

Similarly, Joyce argues that "[u]nderstanding how moral thinking enhanced the reproductive fitness of our ancestors is likely to reveal something about how it continues to be generally useful to us"(Joyce 2006, p.107). Finally, as the quotations from Sinclair and Prinz illustrate, the inference is often implicit and unarticulated, since claims about current function and claims about evolutionary function are often woven together fairly seamlessly.

To be clear, claims about morality's current function can be supported in various ways, but in this paper I will critically evaluate the inference from morality's deep explanatory history to its present function. This genealogical inference should be of independent interest, since it appears to be sound in many domains of inquiry, and since it seems intuitive that genealogical inquiry can help us to learn about the actual function of various social practices. But what does this inference look like in the *moral* case, and how might it be supported?

3 FUNCTIONALIST GENEALOGY

Generally speaking, *genealogy* is a philosophical method which seeks to analyze or evaluate beliefs, values and social practices by studying their etiology.[4] A functional genealogy delves

---

[3] Interestingly, Richard Joyce claims that moral judgment is itself an essentially linguistic phenomenon. See (Joyce 2006, p.111).

into history in order to establish that some value, belief or practice plays a particular role in a social or a psychological system. It may not be immediately easy to see how an entity's history can reveal its function, but, in fact, many philosophers have become attracted to this very idea.

In order to see how this works, we should being by distinguishing between three distinct ideas that might be implicit in the concept *function*. The first is dispositional: the function of an entity is might be the characteristic effect it is disposed to produce within a larger system. The second is etiological: the function of an entity is revealed by the best explanation for its emergence and/or persistence. The third is normative: an entity's function is in part the thing it is *supposed* to do, in whatever sense we choose to attach to that word (Copp 1995, p.110-111). It is important to see that these three ideas can, in principle, come apart. A philosopher might argue that an entity produces characteristic effects within a system without claiming that this is connected to any etiological story, and without claiming that it is 'supposed' to produce these effects.

A functional genealogist is someone who places some kind of priority on the second, etiological sense. The basic idea is this: often, the best explanation for an entity's emergence and persistence will be that it is disposed to produce some effect. When we have a reasonably complete and informative explanation of this sort, it seems natural to conclude that the explanandum's function is to produce those effects. Thus, we might say that token knives have the function of cutting *because* cutting is the effect for which that type has been selected in the past.

However, there are at least two different ways of cashing out that "because". A reductionist follows Larry Wright in reducing the normative and dispositional senses to the etiological one. Wright argued that an entity's function simply *is* whatever effect it has within a system that explains why it exists in that system (Wright 1976). However, this is a very strong claim, and I will not assume that any of my critical targets in this paper are committed to it. I will generally construe them as adopting an alternative, non-reductionist approach, according to which a certain type of etiology is powerful evidence for functional claims. This preserves the idea of a functional-genealogical *inference* without inviting the various objections and counterexamples that follow on the heels of reductive analyses.

In any case, the core idea is that when an entity has a function, this is ultimately *because* it has a certain etiology. In its basic form, this theory remains popular amongst philosophers of biology, and it is the basic framework under which a functionalist-genealogist seeks to provide historical explanations for morality's emergence and persistence (Millikan 1984;

---

[4] The method is very common, and it cuts across the so-called 'analytic/continental' divide in philosophy. See (Rousseau 1755; Nietzsche 1887; Foucault 1977; Williams 2002; Prinz 2007, ch.6).

Neander 1991; Kitcher 1993)[5]. After all, it is eminently plausible that various ethical beliefs, values and practices exist because they regularly produce some set of effects.

A functional-genealogist of this sort can deploy multiple forms of explanation. Most of us are already familiar with one such mode, namely, explanation that cites descent via natural selection. However, Darwinian selection is not the only mechanism which can form the basis of a functional-genealogical explanation. Other modes of explanation are possible. For example, Hume's genealogy of property-conventions begins with the assumption that property-based moral practices are *not* 'natural' features of human societies. He offers a basically game-theoretical reconstruction of the origin of justice which shows it to be a rational response to social conflict, and he can be read as arguing for a functional conclusion, that the amelioration of this same conflict is the primary function of property-conventions (Hume 1739/1740).[6]

So, whether the mode of explanation is explicitly Darwinian, or whether it emphasizes the culturally guided construction of an artificial virtue, the functional-genealogist argues that a foray into the historical origins of a set of beliefs, values or practices can reveal the effect for which it was selected. She concludes that its current function is to produce that effect. In the case of morality, the idea is that the relevant, function-fixing effect is the enhancement of mutually beneficial social cohesion.

So much for the nature of functional genealogy. I will now begin to build my critical case, by arguing that genealogists of morality very often fail to appreciate how demanding their own method is.

4 THE FUNCTIONALIST'S BURDEN: CONTINUITY

We may begin to illustrate the problem by making a simple observation: functions are deeply sensitive to environmental change. This is because nothing has an *intrinsic* function, that is, a role that it plays irrespective of the properties of larger system of which it is a part. More precisely,

(P)   Given some entity E within some larger system S, and given some functional role that E plays in S, there are conditions—extrinsic to E itself—that must obtain within S in order for E to play its role.

Call these *enabling* conditions.

---

[5] Strictly speaking, Kitcher focuses on the notion of a *design*-explanation, but the result is effectively the same, since explanations which cite a kind of design point to selective pressures which ensure that only entities with certain dispositions survive. For criticism of the etiological model see (Cummins 1975).
[6] See (Owen 2010; Williams 2002) for this interpretation of Hume.

The idea of an enabling condition allows us to get clearer on the nature of functionalist genealogy. For it is precisely these conditions which must also play an *explanatory* role in such a genealogy. To take a simple example, we say that polar bear fur has the function of providing camouflage, but only because the bear's *current* environment is white. This means that the color of the bear's environment is an enabling condition. But the explanation for *why* polar bears are now white also cites the whiteness of their ancestral environment. Brown bears migrated north, we are told, and the color of the Arctic environment ensured that those with white fur would have more offspring (Kurtén 1964). Thus, *that* the enabling condition obtained in the bear's ancestral past explains why the species has white fur today. We can draw conclusions about the present-day function of this fur *because* the relevant enabling conditions are also those that help to explain why the fur was adaptive.

Now, suppose a population of polar bears are (cruelly) relocated to the Amazon rainforest, where those conditions do not obtain. Is the polar bear's fur malfunctioning, has it lost its function, or is it simply failing to perform its function? I take no position on these difficult questions, here. But we must remember that if a thousand generations of polar bears retain their white fur in the Amazon, we are forced by the logic of the etiological account of functions to conclude that the fur has lost its original function. Whatever is explaining the trait's persistence, it is not the effect for which it was previously selected.

Taken together, these thoughts suggest a constraint on any functionalist genealogy of any set of beliefs, evaluative judgments, or social practices, in particular those genealogies which seek to infer conclusions about current functions from facts about ancient history. The functionalist-genealogist of morality, as we have seen, wishes to argue that an evolutionary story can reveal morality's function: in coming to appreciate why it arose and persisted, we learn about the function it plays in our own lives. (P) shows that there is a suppressed premise in this mode of argumentation. Call $S*$ the present-day system in which entity E is said to have the function that it (allegedly) developed in S. The strength of an inference about E's function in S* on the basis of its development in S depends vitally on the following condition:

(Continuity)   Given the set of enabling conditions that obtained in S when E developed, the same set of conditions obtains in S* and in the recent history of S*.

A simple way of summarizing this requirement is to note that we can ask: *could* E be reasonably expected to arise and persist in S* or in the recent history of S*? If not, this principle is not satisfied.

I will try to illustrate the point with another example: suppose that a certain model of car once became very popular because it ran well, and it ran well because of an advanced filter-system that removes impurities from the fuel, improving engine efficiency. We might want

to conclude that the filter *now* has the function of improving engine efficiency. (Continuity) requires us to look at the system as a whole in order to rule out certain possibilities. We might suspect that modern versions of the car run on fuel that has no impurities, for example. To rule out these (and other) possibilities is to establish (Continuity), since we will know that the enabling conditions that obtained in the car during the emergence and persistence of the fuel-filter still obtain in contemporary versions of the car. As I will stress later on, failure of continuity does not *entail* the loss of a function, but it does decisively weaken the strength of the functional-genealogical inference.

Now, we will have to pin down just what it means to say that 'the same' set of conditions obtains here. A polar bear's fur will still provide camouflage in a warm environment that has been painted white, even if the original enabling condition (snow or ice) has vanished. So, this principle shouldn't be taken to suggest that enabling conditions must remain continuous in some deep physical or metaphysical sense. Rather, the enabling conditions themselves must first be captured in functionally relevant terms: for the purposes of the function in question, what matters about the bear's original environment was that it was white. The description under which the enabling conditions help us to make sense of the entity's emergence and persistence must remain constant. This is how we should interpret (Continuity).

The (Continuity) principle is the functionalist's burden, and, as we shall see, it makes serious trouble for anyone who wants to derive claims about morality's current function from its evolutionary history.

5   GENEALOGIES OF MORALITY AND CONTINUITY

In the case of evolutionary genealogies of human morality, (Continuity), I will now argue, is not satisfied. Most evolutionary accounts of the emergence and persistence of morality invoke roughly the same kinds of initial conditions, and argue that a standard evolutionary mechanism—natural, sexual or perhaps group selection—can account for the emergence and persistence of human morality (Ruse 1986; Axelrod 1984; Sober 1992; De Waal 2006). These initial conditions are familiar to anyone who's engaged with the literature: Early humans, we are told, lived in small, genetically isolated groups of 50-150 individuals whose interactions were face-to-face; 'defection' (in the game-theoretic sense) was very dangerous to such groups, because they were routinely confronted by resource scarcity and threats from predators. Moreover, collective survival depended on motivationally effective long-term prudential reasoning, but humans are typically bad at motivating themselves in accordance with such reasoning, in virtue of their preference for short-term gain (Kitcher 2006, p.11, 197, 393; Joyce 2006,p. 40-42). These commonly-cited conditions help to explain why human morality is adaptive enough to evolve via natural selection.

Now, here, we come to the observation which drives my critique: a little reflection reveals that these conditions do not in any way characterize modern Western communities.[7] In fact, these differences are *so* striking that we should be lead to wonder why anyone ever thought that the lives of our hunter-gatherer ancestors were of such importance with regards to the question of morality's current function. Hazarding a guess, I would say that these philosophers may be conflating two distinct ideas: first, that *morality* is innate, and second, that morality's *function* is innate and therefore basically a human universal. This conflation might be what is going on in this passage from Sinclair:

> The seeming ubiquity of the moral habit… and the fact that it develops in all normally brought-up humans with little formal instruction is (defeasible) evidence that [its] function is evolutionary. (Sinclair 2012, p.13)

Unfortunately, while a trait might be innate, its function cannot be. A scientist who knew everything there was to know about the ontogeny of a trait could not predict, on the basis of that information alone, the function that the trait would eventually have in the life of an organism. This question can only be answered by the additional examination of the ecological system of which the trait was (and is) a part, as the example of the polar bear shows.

This is what (Continuity) demands of us. And I believe that it is *very* unlikely that (Continuity) will be satisfied for the standard evolutionary genealogy of morality. If we follow biologists in saying that *niche-construction* is the process by which an organism alters the selective pressures on its species by modifying its external environment, then we should say that we live on the other side of a historical process which might be characterized as niche-construction run rampant (Odling-Smee et al. 2003). Human beings have constructed institutions, forms and practices which are bigger and more complex than any ever seen on earth. This process, as I will now show, has entirely erased the enabling conditions which, according to philosophers like Kitcher and Joyce, explain the emergence of human morality and provide it with its original function.

Of course, we might wonder whether the relevant conditions aren't similar *enough*. After all, at a certain level of description, (Continuity) might seem to obtain. We still live in communities, we still need to get along with one another, and we still face various pressing co-ordination problems. However, as I will now show, such descriptions do not reflect the content of the enabling conditions that philosophers and evolutionary theorists deploy in order to explain the emergence and persistence of human morality. It will soon become very clear that morality could *not* be reasonably expected to evolve within communities like

---

[7] I say modern *Western* community because I want to flag that the answer to this may vary from society to society. It is possible that other extant human societies will come closer to satisfying (Continuity), though probably not much closer.

ours, according to our best evolutionary models. I'll begin with Kitcher's model before proceeding to Joyce's.

5.1   THE SHADOW OF THE FUTURE

In order to see why this is so, consider Kitcher's first evolutionary enabling condition: group size. Almost every game-theoretical model which is said to explain the emergence of human pro-sociality requires a small population of participants. Kitcher explains the emergence of psychological altruism by offering a model which is based on *optional games*, or interactions between agents who can choose when and with whom to interact. He defines *discriminating altruists* as agents who "are prepared to play with any organism that has never defected on them, and, when they play, they always cooperate"(Kitcher 1998, p. 501). His aim is to show that discriminating altruists embody a strategy that is collectively stable, or impervious to invasion by rival, less altruistic strategies.

Kitcher defines *selective defectors* as agents who will play against any creature that has not defected on them and who will always defect. He proves that when the population size is small and members do not migrate often, the discriminating-altruist strategy cannot be overcome by selective defectors, because each interaction has a powerful influence over future interactions (Kitcher 1998). In the literature, this effect is called the 'shadow of the future': players are more willing to co-operate because they become worried about their reputation. Yet, when the population grows, interactions with wholly new partners become more common, selective defection becomes much more advantageous. This is an instance of a more general feature of such models: they predict that larger populations are much more vulnerable to free-riders than smaller ones (Kerr and Bruun 1983). For example, Robert Axelrod's highly influential explanations for the emergence of altruism require that the probability that any two group agents will interact more than once is at least .5, and as high as .994, depending on the outcome in question (Axelrod 1984).

In a band of 50 close-knit human beings who rarely migrate to other communities, such probabilities are almost guaranteed. Conversely, in a complex society consisting of millions of highly mobile individuals, they are next to impossible. In modern societies, the shadow of the future is very small.

Actually, the problem runs deeper than this: it is not even clear what it means to speak of a single human 'group' in the modern western context. Kitcher's solution to this problem is to argue that "causal interaction" now makes it the case that there is just a sole human group inhabiting the entire planet (Kitcher 2011, p.305). None of the game-theoretical models I've cited will explain the emergence of altruism even within a population of one thousand members, so it almost goes without saying that a modern group size of seven billion represents a fairly decisive failure of (Continuity).

Now, even if we set this difficult problem aside and allow that there are several distinct human "groups", modern individuals move from community to community with relative ease, and not only will we never interact with most people in our communities, a large chunk of our actual interactions are with near-strangers that we will probably never meet again. This is one important way in which (Continuity) is not satisfied. The sheer size and fluidity of modern populations makes it the case that morality *could* not have evolved in a world like ours.

5.2   JOYCE AND MORAL JUDGMENT

Joyce's explanation for the emergence of the moral sense is multifaceted. When he needs to explain very basic, primitive altruistic motivations, he follows theorists like Axelrod and Kitcher, who explain the emergence of moral dispositions by citing the game-theoretic stability of reciprocity strategies. However, he insists that we do not have an explanation of *morality* until we know how more sophisticated, cognitively-laden judgments arose (Joyce 2006, p.59).

He argues that judgments of the form 'x is required/forbidden' arose because we are, by nature, bad prudential reasoners. That is to say, a human being without the capacity for moral judgment is very likely to sacrifice an overall long-term gain for a smaller, short-term benefit. Since the benefits which arise from increasing group stability are often opaque, group stability itself is often seriously threatened by this weakness. Joyce's hypothesis is that moral judgment is the cure for this disease. Suppose that some non-cooperative action-type *x* provides short-term individual benefits at the expense of the group, reducing an agent's long-term evolutionary fitness. By recruiting moralized emotions such as guilt, the judgment that x is *forbidden* makes it far less likely that the agent will do *x*, increasing both group stability and the agent's long-term reproductive fitness (Joyce 2006, p.110-116).

This is how moral judgment is adaptive, for Joyce. However, in order to assess whether (Continuity) obtains with respect to his story, we need to know what environmental conditions *made* it adaptive. Unfortunately, Joyce rests easy with claims about moral judgment "advancing reproductive fitness" in virtue of providing long-term benefits, but he does not, to my knowledge, provide evidence that this increase was substantial enough to drive the engines of natural selection (Joyce 2006, p.107). We shall have to fill in the details for him.

Fortunately, we have already encountered two widely-cited initial conditions here which make moral judgments strongly adaptive in Joyce's sense.[8] The first and most obvious is resource-scarcity: the sacrifice of greater long-term gain for a lesser short-term benefit can only reduce an organism's reproductive fitness if it will die (or be severely weakened) by the loss of the long-term gain. The second is group instability: it has to be the case that individuals can significantly disrupt group functioning by acting selfishly, or by free-riding on the group's cooperative ventures (Dunbar 2004; Van Vugt et al. 2007).

Do such conditions exist now? Well, one fascinating feature of contemporary Western societies is that these conditions aren't just erased, they're arguably *inverted*. As Adam Smith famously predicted, certain modern forms of economic exchange make it the case that selfishness can contribute to group stability by provoking mutually beneficial *competition* (not cooperation!) and thereby increasing available public goods. This, arguably, has created correspondingly resilient communities which cannot be seriously disrupted by individual selfishness, since they often *feed* on that selfishness. Moreover, contemporary economies thrive on individuals who discount long-term gain in terms of short-term pleasure. Modern economic growth is not encouraged by the cautious spendthrift, it is driven by people who are willing to spend most of what they earn today rather than save for tomorrow (Danziger 2004; Toossi 2002).

In the hunter-gatherer context, private vices may indeed reliably lead to public disasters. However, as Bernard Mandeville saw, private vices can (and do) make for public benefits in a contemporary context (Mandeville 1714). The result is a *cornucopic* society comprised mainly of relatively secure individuals. By comparison with their ancient ancestors, the average European citizen has access to roughly 50 times as many resources, and as a result neither their health nor their survival depends on their making judgments which strengthen either (a) unselfish dispositions or (b) dispositions to ignore short-term gains in favor of long-term profits (Angus 2006). Taken together, this represents a complete inversion of the conditions under which moral judgments had to evolve, according to Joyce's model.

5.3 CONTINUITY FAILURE

I suspect that examples such as the ones just described could easily be multiplied. However, the general point has, I hope, been made. In the case of human morality, (Continuity) fails to obtain along multiple dimensions, and this makes serious trouble for the standard story, according to which the biological and (early) cultural evolution of morality can provide evidence for claims about its current function. If I am right, the evidentiary relation

---

[8] One might worry that I am putting words in Joyce's mouth, here, but since Joyce does not spell out his initial conditions, I can only speculate. Moreover, if he has different initial conditions in mind, my bet is that *they* will also be notably absent from contemporary societies.

between this moral genealogy and the standard functional conclusion is far weaker than has hitherto been supposed.

What I am calling the 'functionalist's burden' is just a result of certain conditions having to play the two roles, explanatory and enabling, that I claim they must play if a functional-genealogical inference is to be warranted. Notice that if we tried to secure more continuity between modern societies and those in which morality emerged, our genealogical account would lose its explanatory power, since it is not at all clear *why* morality should be expected to evolve in a large-scale society under conditions of relative security and resource-abundance. The mere fact that morality might be on-balance beneficial for those living in such a society is not enough to ground an evolutionary explanation; when people register concerns about the explanatory power of *just-so* stories, this is exactly the point they are making (Gould 1980). We will have rescued (Continuity) by robbing the enabling conditions of their explanatory force.

Conversely, re-instating explanatory force by sticking to the original claim—that morality evolved in small-scale, unstable, resource-depleted societies—merely returns us to the problem at hand, namely, that (Continuity) is no longer satisfied. I conclude that the standard evolutionary account of morality does not give us the kind of information we need to say that morality has a certain function. Until we develop functionalist genealogies of morality which satisfy (Continuity), we remain in the dark with respect to such hypotheses.

6 OBJECTIONS AND REPLIES

Thus far, my argument has been fairly simple. If ascriptions of functions on the basis of etiology are to be epistemically justified, (Continuity) must be satisfied. Yet, in the case of the evolutionary etiology of morality, (Continuity) is not satisfied. Therefore, we cannot ascribe functions to morality on the basis of its evolutionary history. I will now try to clarify and extend the argument by listing and responding to a few objections.

6.1 FUNCTIONAL DEFINITION

First, an opponent might accuse me of a certain interpretive blunder. "Surely," they might claim, "the functional genealogist means to say more than this. In particular, the function of morality is not supposed to be some accidental attachment to some independently specifiable social practice. Rather, the maintenance of social cohesion is a *constitutive* function of morality. 'Morality' just is whatever set of practices maintains social cohesion, so it is not possible for it to lose this function."

It is important to see that this response is unavailable to the functional genealogist. This, recall, is someone who wishes to *uncover* the function of morality by looking at its etiology. But to define morality as whatever it is that maintains social cohesion is to abandon this project, for there can be no derivation of morality's function from its history if we have already decided, in advance, that morality is whatever we find in the world that maintains social cohesion. Indeed, this is an inversion of genealogical reasoning, which identifies a target and *then* asks whether the target's history reveals anything of interest. If we have decided *a priori* that the thing we call 'morality' has a certain function, then genealogical inquiry into its functions is pointless.

6.2   IS (CONTINUITY) TOO STRONG?

I have claimed that in order for (Continuity) to obtain between a present-day and an ancestral system, the relevant enabling conditions must remain constant. However, it might be objected that this requirement is too strong. After all, as Kitcher stresses, certain analogues of the basic conditions that obtained when morality arose are still with us. Might it not still be the case that morality enables us to deal with these kinds of problems, even though they are not nearly as urgent for us as they were for our hunter-gatherer ancestors?

Quite so, and this allows me to emphasize an important feature of my argument. My argument, I want to stress, is purely epistemological, and not conceptual or metaphysical. In other words, (Continuity) does not represent a necessary condition on the instantiation of a function, and I have not meant to suggest that a failure of (Continuity) *entails* that a given functional hypothesis is false. Rather, I want to say that when we are conducting specifically genealogical inquiry into the social function of a given set of beliefs or practices, (Continuity) is our best guide. Inasmuch as we wish to form beliefs about functions on the basis of a functionalist genealogy, we must ensure that something like my version of (Continuity) obtains.

So, if we are committed to the view that moral dispositions and judgments can only be reasonably expected to evolve among groups of roughly 50-150 human beings, groups for whom internal disharmony can mean wholesale extinction, then that evolutionary story is of dubious evidentiary value in securing functional claims about contemporary morality.

6.3   SELF-CREATED VESTIGIALITY

Consider the following objection: "Of *course* modern life is not like life in the Pleistocene, and this is precisely what the genealogy of morality predicts. We now live lives of comparative comfort and safety in large-scale, centrally governed societies, but that is *because of morality*. The evolutionary story is one of hunter-gatherers facing a set of

problems; they solved them, and we no longer face such problems because we are their ancestors."

Doubtless, this could all be true. But it is not friendly to the idea that the present-day function of morality is to reduce social tensions in a mutually beneficial fashion. What is interesting about this sort of example is that we have an entity which, *by* performing its initial function, contributes to its own subsequent vestigiality. Our objector is painting a picture of morality that is basically analogous. Yet, this is not a picture which supports the functional claim I am criticizing in this paper, rather, it decisively highlights one way in which evolutionary genealogies can *fail* to support conclusions about function(s) in modern-day contexts. After all, we have just discovered yet another way in which an entity can have a function that it subsequently loses.

Now, there is, presumably, a historical point at which the relevant enabling conditions ceased to obtain in the modern world, when our societies became very large and comparatively secure. Assume, with our objector, that morality played a key role in bringing this about. I have argued that this looks like a case of vestigiality. However, this isn't quite right. In fact, two hypotheses present themselves:

[Vestige] While morality initially enabled us to overcome the conditions, other institutions and practices now keep the enabling conditions at bay regardless of the existence of morality. Morality is, to this extent, vestigial.

[Safeguard] Ever since, morality has acted as a kind of safeguard, one that is triggered only when the relevant enabling conditions reappear.

In other words, either morality's dispositional profile has largely vanished, or its functional dispositions have become *masked* (Johnston 1992)[9]. It might be thought that only the first hypothesis threatens the consensus I am criticizing. It is important to see, however, that the functional genealogists in question cannot accept *either* of these hypotheses. There is no textual evidence to support the claim that when philosophers or biologists say that the function of morality is to produce mutually beneficial social cohesion, they mean something like: 'were humanity to return to a state of resource scarcity and group instability, living in isolated groups of 50-150 members, morality would produce mutually beneficial social cohesion.' In other words, not only is [Safeguard] entirely consistent with my critique, it quite clearly does not represent the position taken by any of the adherents to the hypothesis I am criticizing.

Yet, morality is still with us, and the needs it was designed to meet are (somehow) being met. Isn't this good reason to think that it is still working to meet those needs? This, too, is not a good inference, for, as a matter of historical fact, many large structural changes have

---

[9] I thank Pekka Väyrynen for this suggestion.

produced a large number of other distinct social practices which could easily be fulfilling these needs all on their own.

First, the rise of wealth and cohesion in certain countries is strongly correlated with dramatic increases in state power. At no point in human history have states been more able to capture and punish individuals who do not conform to the law, and this has surely contributed to social cohesion all on its own, largely by punishing free riders. Second, prosperity and security is, arguably, the result of industrial capitalism and technological advancement. It would not be difficult to claim, as some theorists already have, that the current function of *these* institutions is the production of mutually beneficial social cohesion (Runge 1984).

This is why the following inference is deeply problematic:

> Morality evolved to fulfill certain needs,
> Those needs are now being fulfilled. Therefore,
>
> Morality is now functioning to fulfill those needs.

The inference is neither valid nor inductively strong, given the actual development of institutions and practices which almost certainly serve to fulfill the needs in question.

6.4   MORALIZATION

Admittedly, the preceding paragraphs have been overly simplistic, since the idea that we can cleanly separate such phenomena as economics, politics and morality is surely a theoretician's fantasy. In particular, both politics and economics are deeply *moralized*, where by 'moralization' we simply mean the process by which the beliefs, practices and motivations characteristic of morality attach themselves to political and economic institutions (Rozin 1997). As Nico Stehr argues in a recent study, contemporary markets are indeed infused by moral attitudes and practices (Stehr 2006). Moreover, it is easy to see that allegiance to political parties is a deeply personal matter, one that conditions the moral self-image of a great many modern citizens. It is therefore somewhat misleading for me to suggest that the politics or economic activity could, "all on their own", perform the function(s) that morality was selected to perform.

Yet, why should we assume that the moralization of an institution necessarily *promotes* its characteristic function(s)? Reflection suggests that this can go either way, and the empirical evidence shows that the assumption is highly problematic. Stehr himself concludes that morality has both positive and negative effects on economic functioning (Stehr 2006, p.123-135). With respect to the political sphere, no observer of contemporary American politics can fail to notice the positively crippling effects that moralization has in this

domain. According to a widely-held hypothesis—for which there is a lot empirical evidence—major political change in the United States is currently next to impossible, simply because the civic sphere has become permeated with a conception of politics that can be fairly described as Manichean (Tavits 2007; Wisneski et al. 2011). Plainly, the moralization of an institution can promote or hinder its characteristic functions.

6.5 EXAPTATION

Now, to some, it might seem incredible to suppose that such a deeply-ingrained social institution has no active functional profile whatsoever. It is important to see that the loss of a primary function in the past does not entail the lack of a functional profile in the present. This is due to a phenomenon—familiar to scientists but not much discussed by philosophers—called *exaptation* (Gould and Vrba 1998).[10]

Biologists have long been recognized that a fixed trait can, given sufficient environmental change, acquire new functions while shedding old ones. Similar things could easily occur in the case of social institutions. It is plainly possible that morality's properties have made it suitable for this sort of functional co-option. That is to say, while the idea of morality as a social epiphenomenon might strike many as basically implausible, it may simply be that its characteristic role has changed over time.

I think that a model of social function which ignores exaptation is far less naturalistically respectable than one that allows for it.[11] We may be sure that morality must be doing *something* for us, but the possibility of exaptation means that we cannot infer that it is doing the same things that it did in the hunter-gatherer context.

Notice, further, that my observations concerning the moralization of politics and economics actually support an exaptation hypothesis. Timothy Ryan finds that politicians who publicly reject cross-party compromise and consensus-building by deploying concepts such as *right*, *wrong*, *principle* and *integrity* are significantly more likely to be elected, and this must help to explain why morality continues to attach itself to the political sphere (Ryan 2016). This is a classic case of a selection-explanation, since there is significant (artificial) selection-for moralistic politicians. If Ryan is right, then the logic of functionalist-genealogy might force us to say that morality has been subject to exaptation in a certain respect. It is a human trait which may once have fostered unity and prosperity, but which

---

[10] Jason Garson gives an illustrative example: "plant species of the genus *Dalechampia* probably first used resin secretions as a defense against herbivores; later, they became used as a reward system for pollinators" (Garson 2008). See also (Millikan 1984, p.32).

[11] In fairness to Kitcher, he acknowledges the possibility of what he labels "functional generation" (Kitcher 2011, p.237-242). Yet, so far as I can tell, he simply assumes, *a priori*, that subsequently generated functions must be systematically related to what he calls the *original* function of ethical practice. In other words, he rules out a certain form of exaptation by fiat.

has been co-opted into a new system of purposes, tragically serving to perpetuate a system that thrives on *dis*unity and which positively hinders efforts to improve collective prosperity.

All of this being said, I am not officially committed to any exaptation hypothesis, since more evidence is certainly required. I mention this one only to highlight just one of the many ways in which morality's functional profile might be more complex and historically variable than many philosophers seem to admit.

7 CONCLUSION

In this paper, I have identified a common claim about the basic function of morality in human societies. I have not tried to show that the claim is necessarily mistaken, but I have argued that a certain argument in favor of the claim is not promising. Philosophers who wish to deploy the functional claim in support of a larger philosophical program should no longer rely on 'deep' biological or cultural history in order to establish it.

But all is not lost for the functional-genealogist, since there is plenty of human history left on the table. If she wishes to establish her preferred functional hypothesis, she can draw upon more nuanced methods of deriving functions from social history, methods which can remain informative in a society as fantastically complex as our own. Here, Peter Godfrey-Smith's *modern history* theory of functions might be a good start. Godfrey-Smith writes:

> Biological functions are dispositions or effects a trait has which explain the *recent* maintenance of the trait under natural selection. This is the "modern history" approach to functions. The approach is historical because to ascribe a function is to make a claim about the past, but the relevant past is the *recent* past; modern history rather than ancient. (Godfrey-Smith 1994, p.344)

The model is developed precisely in order to deal with cases of exaptation, which deeply problematize the attribution of functions to entities on the basis of their ancient histories. A genealogist of morality who adopted this model would be required to delve into recent cultural history—not something that contemporary analytic philosophers have been keen to do.[12] But it is a task that may be necessary in order to establish the sorts of claims that philosophers have wanted to make about the function of morality.

---

[12] For notable exceptions to this rule, see (Nichols 2002; Prinz 2007).

Works Cited

Angus, M. (2006). *Development Centre Studies The World Economy Volume 1: A Millennial Perspective and Volume 2: Historical Statistics: Volume 1: A Millennial Perspective and Volume 2: Historical Statistics*: OECD Publishing.

Axelrod, R. (1984). The Evolution of Cooperation. *Basic Books, New York. Econornetrica, 39*, 383-396.

Copp, D. (1995). *Morality, Normativity, and Society*: Oxford University Press Usa.

Cummins, R. C. (1975). Functional analysis. *Journal of Philosophy, 72*(November), 741-764.

Danziger, P. (2004). *Why people buy things they don't need: understanding and predicting consumer behavior*: Dearborn Trade Publishing.

De Waal, F. (2006). Morally evolved: Primate social instincts, human morality, and the rise and fall of 'Veneer Theory'. In S. Macedo, & J. Ober (Eds.), *Primates and Philosophers*: Princeton University Press.

De Waal, F. (2013). *The bonobo and the atheist: In search of humanism among the primates*: WW Norton & Company.

Dunbar, R. I. M. (2004). Gossip in evolutionary perspective. *Review of general psychology, 8*(2), 100.

Foucault, M. (1977). *Discipline and Punish*: Vintage Books.

Garson, J. (2008). Function and teleology. *A Companion to the Philosophy of Biology*, 525-549.

Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs, 28*(3), 344-362.

Gould, S. J. (1980). Sociobiology and the theory of natural selection. *Sociobiology: Beyond nature/nurture*, 257-269.

Gould, S. J., & Vrba, E. S. (1998). Exaptation–A missing term in the science of form. In D. L. Hull, & M. Ruse (Eds.), *The Philosophy of Biology*: Oxford University Press.

Hume, D. (1739/1740). *Treatise on Human Nature*: Oxford University Press.

Johnston, M. (1992). How to speak of the colors. *Philosophical Studies, 68*(3), 221-263.

Joyce, R. (2006). *The Evolution of Morality*: MIT Press.

Joyce, R. (2006). Metaethics and the empirical sciences. *Philosophical Explorations, 9*(1), 133-148.

Joyce, R. (2011). Moral Fictionalism. *Philosophy Now, 82*, 14-17.

Kerr, N. L., & Bruun, S. E. (1983). Dispensability of member effort and group motivation losses: Free-rider effects. *Journal of Personality and Social Psychology, 44*(1), 78.

Kitcher, P. (1993). Function and Design. *Midwest Studies in Philosophy, 18*(1), 379-397.

Kitcher, P. (1998). Psychological altruism, evolutionary origins, and moral rules. *Philosophical Studies, 89*(2-3), 283-316.

Kitcher, P. (2005). Biology and ethics. In D. Copp (Ed.), *The Oxford Handbook of Ethical Theory*: Oxford University Press.

Kitcher, P. (2006). Ethics and evolution. How to get here from there. In S. Macedo, & J. Ober (Eds.), *Primates and Philosophers*: Princeton University Press.

Kitcher, P. (2011). *The Ethical Project*: Harvard University Press.

Kurtén, B. (1964). The Evolution of the Polar Bear, Ursus Maritimus Phipps. *Acta Zoologica Fennica*(108), 1-30.

Mandeville, B. (1714). *The Fable of the Bees; Or, Private Vices, Public Benefits.*

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories* (Vol. 14): MIT Press.

Neander, K. (1991). The teleological notion of 'function'. *Australasian Journal of Philosophy, 69*(4), 454-468.

Nichols, S. (2002). On The Genealogy Of Norms: A Case For The Role Of Emotion In Cultural Evolution. *Philosophy of Science, 69*(2), 234-255.

Nietzsche, F. (1887). *On the Genealogy of Morals: A Polemic*: Oxford University Press.

Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (2003). *Niche construction: the neglected process in evolution*: Princeton University Press.

Owen, D. (2010). Genealogy. In M. Bevir (Ed.), *Encyclopedia of Political Theory* (pp. 549-551): Sage Publications.

Prinz, J. J. (2007). *The Emotional Construction of Morals*: Oxford University Press.

Rousseau, J.-J. (1755). *Discourse on the Origin of Inequality*: Oxford University Press.

Rozin, P. (1997). *Moralization*: Taylor & Frances/Routledge.

Runge, C. F. (1984). Institutions and the free rider: The assurance problem in collective action. *The Journal of Politics, 46*(01), 154-181.

Ruse, M. (1986). Evolutionary ethics: A phoenix arisen. *Zygon, 21*(1), 95-112.

Ryan, T. J. (2016). No compromise: Political consequences of moralized attitudes. *American Journal of Political Science*.

Sinclair, N. (2012). Metaethics, teleosemantics and the function of moral judgements. *Biology and Philosophy, 27*(5), 639-662.

Skyrms, B. (1996). *Evolution of the Social Contract*: Cambridge University Press.

Sober, E. (1992). The evolution of altruism: Correlation, cost, and benefit. *Biology and Philosophy, 7*(2), 177-187.

Stehr, N. (2006). *The moralization of the markets*: Transaction Publishers.

Tavits, M. (2007). Principle vs. pragmatism: Policy shifts and political competition. *American Journal of Political Science, 51*(1), 151-165.

Toossi, M. (2002). Consumer spending: An engine for US job growth. *Monthly Lab. Rev., 125*, 12.

Van Vugt, M., Roberts, G., & Hardy, C. (2007). Competitive altruism: A Theory of Reputation-Based Cooperation in Groups. In *The Oxford Handbook of Evolutionary Psychology* (pp. 531-540): OUP.

Williams, B. (2002). *Truth and Truthfulness: An Essay in Genealogy*: Princeton, New Jersey: Princeton University Press.

Wisneski, D. C., Skitka, L. J., & Morgan, G. S. (2011). Political Moralization in the 2012 Election. *Proposal for questions on the 2012 Evaluations of Government and Society Survey. .*

Wright, L. (1976). *Teleological Explanations: An Etiological Analysis of Goals and Functions*: University of California Press.

Cummins, R. C. (1975). Functional analysis. *Journal of Philosophy, 72*(November), 741-764.

De Waal, F. (2013). *The bonobo and the atheist: In search of humanism among the primates*: WW Norton & Company.

Foucault, M. (1977). *Discipline and Punish*: Vintage Books.

Garson, J. (2008). Function and teleology. *A Companion to the Philosophy of Biology*, 525-549.

Joyce, R. (2006). *The Evolution of Morality*: MIT Press.

Kitcher, P. (2011). *The Ethical Project*: Harvard University Press.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories* (Vol. 14): MIT Press.

Nichols, S. (2002). On The Genealogy Of Norms: A Case For The Role Of Emotion In Cultural Evolution. *Philosophy of Science, 69*(2), 234-255.

Nietzsche, F. (1887). *On the Genealogy of Morals: A Polemic*: Oxford University Press.

Owen, D. (2010). Genealogy. In M. Bevir (Ed.), *Encyclopedia of Political Theory* (pp. 549-551): Sage Publications.

Prinz, J. J. (2007). *The Emotional Construction of Morals*: Oxford University Press.

Rousseau, J.-J. (1755). *Discourse on the Origin of Inequality*: Oxford University Press.

Williams, B. (2002). *Truth and Truthfulness: An Essay in Genealogy*: Princeton, New Jersey: Princeton University Press.