The Evolution of Moral Knowledge

Richmond Campbell

Abstract

That moral knowledge is objective has been relentlessly debunked by reference to the bio-cultural evolution of human morality by natural selection. Building on Darwin's fundamental insight that morality is grounded in animal sociality, I offer reasons to reject the traditional human-centered conception of morality and argue that the evolution of moral knowledge by natural selection predates humans and the core elements of moral knowledge are in a relevant sense objective.

In Sally Rooney's novel Beautiful World, Where Are You, Alice writes to a close friend about knowing the difference between right and wrong.

> … most of our attempts throughout human history to describe the difference between right and wrong have been feeble and cruel and unjust, but … the difference remains – beyond ourselves, beyond each specific culture, beyond every individual person who has ever lived or died. And we spend our lives trying to know that difference and to live by it, trying to love other people instead of hating them, and there is nothing else that matters on the earth.[1]

Alice is a fictional character, but her sentiments reflect a moral perspective that many of us share. Unfortunately, it faces bedeviling puzzles if morality evolved through natural selection, as many of us now believe. I offer a different way of thinking about these puzzles by drawing on Darwin's idea of the continuity of moral evolution from animals to humans. I argue that the puzzles can be resolved if we reject the usual human-centered conception of moral knowledge.

Three Puzzles

Transcendence?[2] If the truth about right and wrong is objective, transcending individuals, including their beliefs, emotions, and evaluative attitudes, indeed, transcending all cultures, then it appears that the explanation of why

morality exists cannot be dependent on how our moral beliefs, emotions, and evaluative attitudes evolved by natural selection. The two ideas about morality—transcendence beyond the contingences of evolution and dependence on those contingencies—seem manifestly incompatible.[3]

Pluralism, not universality. Empirical studies of moral norms across cultures show that despite great variance certain categories of moral norms appear in virtually all cultures: those forbidding harm to others, disloyalty to family and community, betrayal of trust, interference with the equal liberty of others, and being unfair. It is plausible that these norms evolved by natural selection.[4] That said, none expresses a universal moral truth about the difference between right and wrong that takes moral priority over the other norms. Instead, we see a pervasive moral pluralism regarding their priority, sometimes relativistic,[5] that is contrary to the sentiment that Alice seems to express.[6]

Personal Morality?[7] Some aspects of morality are obviously personal and relevant to loving others. We care about loyalty to family and community, for example, but not to complete strangers. But morality that evolved in large groups composed mostly of strangers is not personal in this way.[8] We ought not to harm those we love, but we also ought not to harm strangers. Our obligation to keep promises or to be fair or to respect the equal liberty of others also applies to strangers. Given how much our capacity to survive in large groups has depended on the moral motivation of strangers, natural selection must favor more than a capacity not to harm those we love and to be loyal to them.

Though these puzzles are far from trivial, I will argue that Alice can resolve them if she rejects the standard human centered perspective in which morality did not exist before humans. While Darwin did not reject this perspective in The Descent of Man, he allowed that "…the difference in mind between man and the higher animals, great as it is, is certainly one of degree and not of kind."[9] He held, more exactly, that the capacity for morality in higher animals resides in their sociality, despite striking differences in the degree of intellectual powers.[10] My plan is to build on his insight to suggest that moral knowledge evolved by natural selection prior to humans and its most basic elements are unchanged.

Three caveats are in order before I begin. First, I will not compare my approach to important alternative attempts to defend moral knowledge based on a human-centered conception of morality in which moral beliefs can be objectively true despite evolutionary influences.[11] That task, though philosophically familiar, is arguably redundant if the animal-centered approach is shown to resolve the three puzzles just presented since the

puzzles would be resolved for human animals as well. Of course, it may be doubted that an animal-centered approach is genuinely about morality. That concern takes us to the next caveat.

Second, I will assume that many animals other than humans can be motivated to follow moral norms and do so without reflection on whether their motives are moral. This assumption runs contrary to the philosophies of Aristotle, Kant, and Hume.[12] On the dominant view in Western philosophy, for animals to be morally motivated, they must be capable of reflecting on the moral status of their motives and animals without language are incapable of the meta-cognition needed for this self-reflective task thought to be necessary for moral choice. While I concede that lack of sophisticated meta-cognition entails an absence of sophisticated moral agency, I concur with Mark Rowlands among others that animals lacking in sophisticated agency can still be moral.[13] Animals are moral when, for example, they are motivated to defend their group in the face of danger because they have internalized a norm to do so and the norm functions morally, even if they lack the capacity individually and collectively to reflect on the moral nature of their motivation. But what is it for a nonhuman animal to internalize a norm and what exactly would make it moral in the absence of self-reflection? Moreover, how is moral knowledge possible without self-reflection? These questions take me to the third caveat.

This final caveat is that I will not try to answer these questions by appealing to a priori principles. Instead, I base my answers on an evolutionary concept of moral normativity and moral know-how that can explain the gradual evolution of moral knowledge that began before the evolution of humans.[14] I shall assume, following Sarah Vincent, Rebecca King, and Kristin Andrews,[15] that to be moral it is enough for animals to be motivated to conform to moral norms. My focus will be first on normativity alone, then on moral normativity, and finally on moral knowledge to explain how animals other than humans can know morally how they should act. Key to understanding animal moral knowledge is to understand the "evolutionary function" of moral norms in social animals.[16] Since moral knowledge so conceived is dependent on motivation, it may appear to be stance dependent and not open to a moral realist interpretation in which moral knowledge is stance independent. I will address this complex issue in due course.[17]

My main aim, consistent with these caveats, is to show how an animal-centered approach to moral knowledge can in principle dissolve the above three challenges while not importing any assumptions not fully compatible with a thoroughgoing Darwinian account of moral evolution.

On the Origin of Moral Normativity

I begin by drawing on the concept of "naïve normativity".[18] To be normative in this basic way, animals must associate themselves with a group of animals that they readily distinguish from other animals. Second, they must be able to tell when behavior in their group is voluntary, such as one animal choosing to push another instead of being pushed into another by the movement of other animals. The third condition is more complex but still easily satisfied by animals other than humans. The animals must be able to recognize certain patterns in their group's voluntary behavior, such as forming a circle in moments of perceived external threat,[19] and they must be motivated to conform to the pattern in part because the others do. The fourth condition complements the third: that animals in their group are motivated to respond hostilely to group members that do not conform to the pattern. I will assume that normativity of this naïve kind is not in question, however challenging it is to understand its causal complexity.[20]

Kristin Andrews provides a poignant example drawn from her research on chimpanzees.[21] One group discovered a more efficient way to crack open a nut than is practiced in most groups. When a few females among them migrated to a group ignorant of the easier method, the newcomers chose to follow the established pattern of their new group. Why not try to teach the easier method to the new group, if not the recognition of the strength of the current norm and the risk of being socially rejected for not conforming?

Not every norm is a moral norm, and certainly we do not perceive norms about breaking open nuts as moral norms. Indeed, as amply illustrated by Sarah Vincent and her colleagues, normativity is broader in scope than moral normativity.[22] What then is the difference?

The difference lies in how moral norms function. Note that the concept of function has a familiar interpretation in evolutionary terms. The function of the heart is to pump blood to deliver oxygen to various parts of the body. That is why the heart exists, but this reason can be explained through the mechanism of natural selection. This trait of the heart, to pump blood with oxygen to the body, was selected for, among other variants, because it increased the reproductive fitness of animals whose hearts have this trait. In the literature this is known as the "selected effect" view of function.[23] It provides a causal explanation of why something exists by reference to the theory of natural selection. Thus, eyes function to enable animals with eyes to see. That is why they exist, if we understand this teleological talk in Darwinian causal terms.

What is the moral function of norms? I believe that this question is answered using the selected effect understanding of function. The role played

by moral norms in virtually all known human societies is to facilitate mutually beneficial sociality. Think of moral norms for helping others avoid harm, being loyal to family and friends, keeping promises, respecting the freedom of others, and being fair in allocating resources. While these norms become qualified in institutionally structured modern societies where they are juxtaposed with norms of authority and purity, they all tend to make interdependent living more mutually beneficial than it would otherwise be.[24] We can say in sum that the function of moral norms, taken individually as well as together, is to facilitate interdependent living. The specifics of moral norms can vary in interpretation across cultures and continue to evolve,[25] but their function, as selected effect, is always to make interdependent living among social animals more successful than it would be without those norms. That is why moral norms exist.

The function of moral norms is unique in the deep way they resolve problems of interdependent living. Conventional norms also improve interdependent living. Norms of greeting, for example, make recipients feel welcome and put them at ease. They are not moral norms, however, and failure to follow them need not be a moral failure. Moral norms, by contrast, are deeply bound up with survival, as in the examples of avoiding harm, being loyal to family, friends, and community, reciprocating favors and being trustworthy, respecting the liberty of others, and treating them fairly. For this reason, adherence to them tends to be motivated by moral emotions of sympathy, loyalty, trust, and respect, as well as the reactive emotions of resentment and guilt, and to take priority over norms of convention when moral norms conflict with them.[26]

That said, it is important not to confuse the function of moral norms with the motivation to conform to them. When animals conform to norms, they are moved to follow the patterns of behavior voluntarily chosen by associates, say about how to crack open a nut. The case is similar when the norm is moral, but the motive is stronger and takes precedence in cases of normative conflict. If the moral norm is for stronger adults in a group to form a protective circle around the vulnerable when danger is imminent, the adults have learned to do so from the strongly motivated behavior of other adults in response to perceived danger. Others are moved even when the danger is great, and the strength of their motivation is contagious and resilient in the face of distractions. Younger adults want to follow the example of the older adults, and the older ones will have learned from adults before them. While the result can be vital to their success at interdependent living, this function of their normative behavior will not be their motive for conforming to the

norm. They conform to the norm because following it has become inherently attractive to them and that in turn improves its function.

Compare human motivation to conform to modern rules of the road when driving. Various motives are relevant, such as to avoid danger or legal penalties, but we follow the laws of the road, in large part, because others do so and expect us to do the same. Our motivation can be reinforced by moral emotions, as in feeling sympathy when we avoid harming another or in feeling respect when we refrain from interfering with someone's liberty without cause. Still, moral norms also move us just in themselves even as they function socially to deeply enhance interdependent living and thereby improve our overall chances of survival. Moral norms have evolved to move us in this direct way for this reason.

The application of moral norms is also guided by moral consistency reasoning, though in the case of most animals, including some humans, the reasoning need not be self-conscious. If a norm calls for protection of the vulnerable when danger exists, a member of the group who does not respond appropriately acts inconsistently with the norm and is subject to hostility from others in the group. Inconsistency in moral response can be recognized by the group with the same accuracy that they recognized danger, but without the meta-judgment that they are responding to moral inconsistency. Reflection on moral inconsistency can be critical to successful moral reasoning but it need not be. What is necessary is a capacity to distinguish behavior that fits the relevant norm from behavior that is contrary to it. In the latter case, an animal motivated to conform to the moral norm will avoid behavior that is inconsistent with it and react negatively to such behavior in others. Though extremely basic, this capacity for moral reasoning is important for moral norms to function properly in a group.

In sum, four elements are normal ingredients of moral normativity in a group. First, its members display at least naïve normativity with respect to the norm in question. Second, in virtue of this normativity, the members can live together interdependently in ways that augment their mutual survival. Third, they are normally motivated to do so by moral emotions, like sympathy and loyalty, that are integral to their experience of moral norms. Finally, they are guided by a capacity to see when behavior would or would not fit the norm in question and respond negatively to behavior that does not.

The Is-Ought Gap

This understanding of moral normativity seems to commit the naturalistic fallacy by inferring 'ought' from 'is'. It does not, but the reason may not be obvious. I do hold that the values inherent in successful interdependent

living, like staying healthy, being mostly safe from external harm, having healthy offspring, and enjoying life, are explained themselves by reference to natural selection. Do I not, then, imply that all moral value can be explained by reference to natural selection and thereby commit the naturalistic fallacy?

Notice that if I were to reject the last claim and take the normative stance that these aspects of successful living are good independently of how they evolved, I would face a dilemma. The normativity contained in successful living cannot be itself moral, given my story, for then my account of moral normativity would be circular. On the other hand, if these values contained in successful living are not moral values, what makes the function of group norms moral when they result in the non-moral value of contributing to successful interdependent living? My answer cannot be that this Darwinian evolutionary function is moral simply by being this kind of Darwinian evolutionary function, for then I would infer a moral 'ought' from an evolutionary 'is'—I would infer that the function of moral norms is moral just because they have the evolutionary function described. Indeed, it may appear that I have no other alternative if I continue to insist on explaining moral normativity by reference to Darwinian evolution.

My reply follows. Drawing on Anscombe's discussion of 'is' and 'ought',[27] we can distinguish two ways that representations of the world can function, whether they are linguistic or not. They can function to represent the world as it is, or they can function to represent the world as it ought to be. In the first case, if you find that the world doesn't fit the representation, you revise how you represent the world; in the second case, when there is a mismatch, you try to change the world. The first kind of function is descriptive and the second normative. Obviously, from the fact that a representation functions descriptively, it doesn't necessarily follow that it functions normatively. Notice, however, that it is possible to represent accurately how a norm functions normatively without using this meta-normative representation normatively. The latter is exactly what I am attempting to do in this paper. Of course, I may also approve of how the norm functions and in this regard take a normative stance toward it. I may hope that you do as well. Yet in describing the normative function, I need not thereby use my descriptive representation of normative function normatively. For this reason, I do not commit the naturalistic fallacy.

This brief reply to the charge of committing the naturalistic fallacy needs to be unpacked and qualified to be deemed satisfactory. The qualifications and the unpacking will be offered in the sections to follow that provide an animal-based conception of moral knowledge and lead to a stable resolution of Alice's puzzles.

Two preliminary points are in order before we move on. First, a representation can serve both functions at once. For example, I can describe persons as being in pain, with the normative implication that, other things equal, they ought not to be this way. Or a hen can call her chicks, simultaneously announcing the location of food and bringing them to her.[28] In both examples descriptive and normative functions are exemplified at once. Thus, the distinction between descriptive and normative functions is compatible with a moral naturalism where a fact about the world functions as both descriptive of the world and morally normative for the world. We will return to this point in due course.

Second, even when the normative function of norms is moral, moral norms may fail to function properly when they are a product of natural selection. Some hearts and some eyes function very poorly. Even when their evolutionary function is to pump oxygenated blood or to see, hearts and eyes may do so poorly or not at all. The same holds for norms that function morally. We can say that they ought to function differently than they do, even when their function is determined by natural selection. Notoriously, some moral systems are morally corrupt and unjustified,[29] but in a generic sense they are still moral systems and people can mistakenly treat them as having normative status when they ought not to have it. Hence, once again, my account of how moral norms evolve does not thereby commit me to a normative stance about what evolves.

On the Origin of Moral Knowledge
The presence of moral normativity does not ensure moral knowledge. We need to explore what is missing. To do this, we need to bring into focus elements of moral normativity that lie beneath the surface when moral norms function as they are designed by natural selection. They were touched on earlier but need elaboration.

When an animal voluntarily chooses to follow a pattern that functions as a moral norm, the pattern contains what Jennifer Woodrow and I have called reason-response pairs.[30] If the norm is to protect vulnerable individuals threatened by danger, say by making a protective circle around them,[31] the protective behavior is triggered when group members recognize that vulnerable members are in danger. In the context of naïve normativity, this recognition is the "reason" for protecting them and protecting them is "the normative response" to this reason. The animals who respond need not have any concept of a reason. Still, the need that prompts protection functions as a reason, since the need is part of a reason-response pattern that functions as a moral norm. In effect, the representation of vulnerable individuals needing

protection functions normatively to move others to change the circumstances so that those threatened are protected. Individuals with a reason to protect normally respond consistently with the response pattern of the others in the group. When some respond inconsistently, say by choosing not to protect when there is reason to protect, other members of the group respond negatively to this inconsistency. To those others, the inconsistency functions as a reason to register their displeasure.

The representation of the reason for protection functions, in part, descriptively and can be true or false. If there is in fact no danger when signs of danger are perceived as a reason to protect, there is no objective reason to protect. The representation of danger fails to mirror reality. Similarly, if others misperceive consistency in response to real or apparent danger as inconsistent, here too there would be no objective reason to correct the individual who appears to respond to danger in ways inconsistent with the norm. In both kinds of case, a distinction exists between the subjective element, namely the perceived reason, and the reality represented. In both cases, the normative response, to protect the vulnerable or correct those acting inconsistently with the norm, may not be justified from an objective standpoint, and thus not reflect moral knowledge.

In general, for animals that respond to moral norms, two necessary and jointly sufficient conditions exist for an animal to act with moral knowledge. First, the normative response must be motivated by a reason that reflects how the world is from an objective standpoint. The perceived danger or inconsistency, to use our examples, must be real.

The second condition for moral knowledge is more complex. It concerns the normative response to a reason. The response in our example is acting to protect the vulnerable. In this case there is a representation of the world, namely one in which the less vulnerable make a protective circle around the more vulnerable. This time, however, the representation functions in a normative fashion (in Anscombe's sense) to indicate how the world should be changed in response to the reason to protect. As before, there is a subjective element and an objective element. The subjective element is the representation of how the world should be, given the reason to act, together with the strong motivation to help create the represented change. Before we move to the objective element, let us linger a bit on the subjective.

It is tempting to ask why belief is not part of this subjective element. Of course, non-human animals do not have beliefs in the form of linguistically formulated propositions about moral truth. But arguably, moral beliefs are not essential for moral knowledge even in the human case. Imagine a highly competent woman passed over for job promotion in favor of a less competent

man. She may know emotionally that she has been treated unfairly (based on moral consistency reasoning that is emotion-based and spontaneous) even though for ideological reasons she believes the decision was fair.[32] As in this human case, animals need not have moral beliefs to have moral knowledge.

But what is the corresponding objective element that would justify the normative response? To answer this question, we need to step back and consider the nature of the animal moral knowledge at stake.

As indicated earlier, I am assuming moral pluralism for humans and other animals, at least in the sense that there is no single moral norm that determines what is morally right in every situation. More than one norm exists for both humans and other animals. That means: the objective normative element of moral knowledge cannot reside in general in understanding what is required by any one moral norm, such as some form of utilitarianism or contractarianism. Rather, moral knowledge of what to do is having the moral know-how to respond with moral consistency to the moral norms and the reasons to act presented in the situation at hand when the norms are fulfilling their moral function and the reasons are accurate.[33]

This condition is complex, in part because moral know-how is manifested in moral consistency, that is, in treating like cases alike. To circle with others in response to danger in one case but not in another is morally inconsistent unless the two cases differ in some morally relevant way. Suppose that an adult animal conforms to the norm of caring for an ailing offspring and is thus moved to do so by feelings of sympathy and loyalty not to abandon the offspring to make the circle. Whatever choice the adult makes will be morally inconsistent with one of the two norms unless, as is likely, the group treats caring for a sick offspring as a morally relevant difference.[34] Complexity can arise also from uncertainty about whether the norm in question is functioning morally, as when too few are following the norm for it to improve interdependent living. There are many cases, however, where clearly only one norm applies (say, the norm to protect the vulnerable, as in our first example), the reason to protect (the presence of danger) is clear, and the system of moral norms functions well enough for animals in the group to do well.[35] In these many straightforward instances, animals that are not human display the relevant kind of moral know-how to have moral knowledge. They know morally what to do.

As indicated early on, I do not assume that nonhuman animals have the meta-cognition necessary for language and sophisticated moral agency that can allow them to reflect on the moral status of their motives and form propositional moral knowledge, arguably manifested in knowledge of the Golden Rule.[36] Most cases of purported human moral knowledge, such as of

the Ten Commandments, do not require recursive thinking beyond the cognitive powers of nonhuman animals. It is possible, therefore, to concede that some features of human moral cognition are beyond the powers of other animals without supposing that human moral knowledge in general is different from that of animals, indeed so different as to constitute an altogether different kind of moral knowledge.

In the previous examples of animal moral knowledge, the less vulnerable protect the more vulnerable out feelings of sympathy and loyalty. There are, however, examples of animal moral knowledge where these feelings are secondary. When a young female elephant neglects her off-spring, allowing to it lag behind family and possibly become lost or injured, she is apt to incur the wrath of older females, even if unrelated to her, because the young mother is acting wrongly. Carl Safina reports an incident when Tecla, a familiar female adult, complained loudly when the truck that he was in had come between Tecla and an elephant baby.[37] Those with him in the truck took Tecla to be complaining at them, but it was soon evident that her concern was directed at the baby's mother who was nearby yet not attentive to the baby's danger. Safina calls this situation "understanding third party relationships" since Tecla is a third party in relation to the mother and her offspring. Tecla understood that relationship and, arguably, was reinforcing the moral norm that applied to it, not acting out of sympathy for a calf that was not hers. The calf's mother cared about her calf but not sufficiently to be mindful of possible danger in this instance and needed to be reminded to protect her calf, something she knew already, indeed, knew morally, without language.[38]

Are there examples of moral knowledge among nonhuman animals that not only do not depend on feelings of sympathy and loyalty, but also do not relate directly to the core moral norms listed earlier? Consider the role of an older female elephant who has the experience and memories to serve as the matriarch of her family and those bonded to it. Her leading others in life-and-death decisions, such as leading them in a drought over a period of days to find food and water, functions as a moral norm that requires the others to follow her lead whether they feel sympathy or loyalty or neither. It is her status as the female leader that calls for all those in her extended family to follow her to survive. The followers cannot know enough from experience to assess her competence. They follow, not from prudence, but because they know morally what they must do.[39] Their interdependent adherence to the moral norm of following the matriarch serves their survival more often not.

Obviously, there is no exact point in time when moral knowledge first appeared, since the necessary and jointly sufficient conditions for moral

knowledge are not precise. Still, they were fulfilled at some point much earlier than the emergence of the human genus. Elephants are one example among countless others.

It might be objected that the jointly sufficient conditions for moral knowledge are open to counterexample if the relevant group of moral equals is restricted, say where some non-Africans claim falsely to know that treating Africans as chattel slaves is morally permissible, some males claim falsely to know that it is morally permissible for them to treat females as subordinate, or some humans claim falsely to know that it is morally permissible to treat (other) animals inhumanely. Are not both jointly sufficient conditions met in such cases? Good arguments in the literature imply that neither is. The boundary between equal moral inclusion and moral exclusion would in each case misrepresent morally relevant factual similarities between groups and the morally consistent application of norms.[40]

Is the Case for Moral Knowledge Normative?

Before we return to Alice's sentiments, it is well to take stock. I intend that the foregoing story is descriptive rather than normative. I have described what naïve moral normativity and moral knowledge are among animals who have evolved the social capacity to be moral. The two jointly sufficient and necessary conditions for moral knowledge are descriptive, and nothing normative is intended to follow logically from them (on pain of my committing the naturalistic fallacy). I have claimed that moral normativity and moral knowledge evolved in animals before humans, but I have not claimed that the world ought to be this way. The basic elements of moral normativity and moral knowledge can be objectively and universally true of moral evolution even if neither Alice nor I have any normative commitments in this regard. Of course, we do share normative commitments, for example, to know what is right and wrong, but they are not needed to understand facts about moral evolution.

It may be objected, however, that facts about moral knowledge are intrinsically normative since moral knowledge requires justification for its representation of how the world ought to be and justification is inherently normative. Is there a satisfactory reply? My observations about how moral knowledge evolves are meta-ethical and are devoid of any normative import about what is justified, just as my observations about when norms have a moral function are not offered to justify the function of moral norms. When animals without language follow the moral norm of making a circle to protect the vulnerable, they are justified with respect to that normative activity when danger is present and when they act with moral consistency in relation to

norms that function as moral norms. In saying as much, I imply nothing about how they ought to act; I only describe what is morally normative for them. While I am committed with Alice to knowing in our own lives the difference between right and wrong and acting on this knowledge, my factual claims about moral evolution are logically independent of that commitment.

Suppose I say to someone in all seriousness: "I know what you are doing is wrong and you do too."[41] Am I not expressing a normative commitment in opposition to the actions of the person addressed? I agree that I am, but I would not be making a meta-ethical statement about how moral knowledge evolves. My position here is that such a meta-ethical statement would carry no normative commitment.

It may also be objected that the jointly sufficient and necessary conditions for moral knowledge, though descriptive, are not "stance independent" in the sense used in the literature, for example, by Sharon Street,[42] since they refer to the states of mind of the animals that would have moral knowledge. My reply follows. When I claim in the next section that moral knowledge is "objective and universal" I mean to imply only that the core conditions for moral knowledge (as it evolved prior to the Homo genus) apply to all animals with the capacity to fulfill them. This general meta-ethical claim about the conditions needed for moral knowledge to evolve is itself "stance independent" in that its truth would not depend on anyone's beliefs, desires, or feelings. I believe that this understanding of moral knowledge is in the spirit of Alice's sentiments stated at the beginning.

It should be clear that the necessary and jointly sufficient conditions for moral knowledge do not themselves constitute a moral norm. The reason is that they do not function as a norm that motivates animals to live together interdependently. Rather they describe the conditions under which animals could follow norms that would have this result. These conditions are not themselves norms that specify actions to be performed when certain triggering conditions are evident, such as the action to form a protective circle around more vulnerable members of the group when danger to the group is present. The meta-normative conditions for moral knowledge are simply too general to function in the specific ways that moral norms function among animals.[43] This fact is not a problem for Alice if she is referring to the meta-normative conditions for moral knowledge when she talks about the difference between right and wrong being objective and universal.

Meta-Normative Conditions for Moral Knowledge
The meta-normative conditions that give the evolutionary origin of moral knowledge are of three kinds. There are the conditions for moral normativity

(described four sections ago) that promote interdependent living needed for survival. Second, there are the conditions necessary and jointly sufficient for the normativity that exemplifies moral knowledge (described two sections ago). These two sets of meta-normative conditions evolve by Darwinian bio-cultural selection—the third main meta-normative condition. A full meta-normative exposition of how moral normativity and moral knowledge evolved is beyond the scope of this essay but not necessary to explain how Alice can address the puzzles that she faces.

Consider these three main meta-normative conditions for moral knowledge in reverse order. I take the explanatory scope of bio-cultural evolution by natural selection for granted, though the legitimacy of Darwinian natural selection as a mode of explanation has been challenged on multiple grounds, for example, as an untestable hypothesis or a hypothesis that is either limited in scope or empirically false.[44] These objections among others, however, have been carefully examined and rebutted in the literature.[45] Still, we need to consider whether Alice would regard bio-cultural evolution as appropriately objective. Since it is about cultural evolution, it is in that sense not independent of culture. On the other hand, notice that no specific culture explains the existence of bio-cultural evolution. Alice can, therefore, say (as noted) that the difference between right and wrong is "beyond ourselves, beyond each specific culture, beyond every individual person," even when that difference is determined by bio-cultural evolution. Meta-normative bio-cultural evolution is not a product of any culture or person; hence, in this respect at least, the difference between right and wrong would be objective.

What about the second main meta-normative condition for moral knowledge. That consists in two necessary and jointly sufficient conditions for normativity that exemplifies moral knowledge in at least the sense of moral know-how. The first is that the relevant normative response (such as forming a protective circle around the vulnerable) must be motivated by a reason that reflects how the world is objectively (such as there being real danger to the vulnerable). Whether this condition is satisfied is clearly an objective matter, independent of emotion and belief. The second condition, however, requires a morally consistent response to the relevant moral norms given in the situation at hand. Is this second condition objective as well? There are two parts to this question. First, is it an objective matter what are the relevant moral norms? Second, is the achievement of moral consistency in response to moral norms an objective matter?

Although core moral norms, as we noted earlier, vary in interpretation across cultures, all cultures include these core norms.[46] They are not the

product of any culture, much less an individual person, but are the product of bio-cultural Darwinian evolution and have been sustained in diverse environments. Since their existence is not dependent on any persons or culture or on similar environments, these moral norms are objective in the sense Alice intends. Moral norms among non-human animals are more primitive, such as the norm to form protective circles to protect the vulnerable, and fewer in kind, but they are no less objective for being primitive and fewer in kind.

What about moral consistency? The problem of moral inconsistency arises for all animals with moral norms. Recall the example of a mother choosing to tend her sick calf instead of helping to form a protective circle in the face of danger. We imagined that tending a sick calf is a morally relevant difference for her herd, so that there is no moral inconsistency in the mother's choice. But another group might not view caring for a sick calf as a relevant difference, making her moral concern for her calf morally inconsistent with the norm to help to form a protective circle in response to danger. It would be tempting, therefore, to see moral inconsistency as stance dependent and not objective. That would be a mistake. The meta-normative requirement to maintain moral consistency arises from the fact that inconsistency in responses to moral norms directly undermines the capacity of moral norms to guide group behavior. Once moral inconsistency is tolerated, it becomes unclear what the moral norms require or indeed whether they are moral norms since they would to some degree cease to promote group survival.[47] While the perception of moral inconsistency can vary, the need for moral consistency is independent of culture and the individuals who embody culture.

Finally, there is the matter of what makes norms moral. Is this an objective matter? Norms are moral when they systematically promote survival of a group through successful interdependent living.[48] Cultures that include such norms tend to survive better than otherwise, but the cultures are moral in virtue of the work done by these norms. The norms are not moral in virtue of the cultures in which they are found. Nor are the norms moral in virtue of the perspectives of persons within the cultures. What makes them moral is objective, namely their causal effect on group survival. In sum, the three key ingredients that constitute the evolution of moral knowledge are fully objective in the sense Alice expresses when she implies that there is an objective difference between right and wrong.

Even if we grant when Alice speaks of the difference between right and wrong transcending culture, she is talking about these three underlying evolutionary conditions that allow specific moral norms to evolve, we must

allow that she values the existence of these conditions. It matters to her that they exist. I grant this important point. Because she cares about right and wrong, it matters to her that there is an objective and universal underlying explanation of the difference between right and wrong. It matters to me too. But the evolutionary conditions that explain the difference are not themselves normative; nor is the description of them normative. They are, just in themselves, norm independent facts about moral evolution.

Alice's Sentiments Undebunked

The puzzles were that moral evolution by natural selection is incompatible with Alice's apparent sentiments about moral transcendence and universalism and the moral importance of love. But if the foregoing account of how moral knowledge evolved is anywhere near the truth, moral knowledge requires reasons for action located in fact, responses to those reasons that are guided by norms that function morally, and moral consistency in the responses to those reasons—summed up in three conditions for the evolution of moral knowledge in the previous section. Does this general understanding of moral knowledge resolve the puzzles?

To see how it does, we must first recognize an ambiguity in the first two puzzles. Take the first puzzle. In one sense, the summary entails that moral knowledge is not transcendent of time, place, and culture, but depends on how moral norms evolve and how they are applied to what is happening on the ground. Without such contingent facts moral knowledge could not exist. But this observation is consistent with the transcendence of moral knowledge in a second sense. The three meta-normative conditions explaining the evolution of moral knowledge hold no matter what our stance towards them might be if they are as I describe them. If that is what Alice could mean when she says that the difference between right and wrong is beyond individual persons and each specific culture, then the puzzle about transcendence dissolves. The difference between right and wrong would in one fundamental sense transcend changes in our beliefs and feelings and our surrounding culture.

As for universal moral norms, Alice doesn't speak of a single moral norm that determines all moral truth. She talks instead about trying always to discover "the difference between right and wrong." That difference can be the same in the sense just given yet vary in detail from situation to situation and be compatible with moderate moral pluralism in which no single norm determines in all cases what is right or wrong. The first two puzzles, in short, present no real problem for Alice's perspective if we understand (descriptively) the ancient roots of moral knowledge that evolved long before

there were humans. In one sense there is no transcendent universal moral knowledge. In another sense moral knowledge is rooted meta-normative facts that are both transcendent and universal.

Does moral objectivity so conceived entail realist moral naturalism? It does in at least one robust sense of the term, since there would exist natural facts about which we can be mistaken that determine whether an action is morally right or morally wrong.[49] In our example of animals protecting the vulnerable in the face of danger, what they do is the morally right thing to do in virtue of the natural facts of the situation as spelled out above. Arguably, animals do not have the propositional knowledge that what they do is morally right, but what makes it morally right is entirely contained in the natural facts of the case, independently of our opinion on this matter. Nevertheless, there is a stronger interpretation of moral realism that goes beyond this understanding because it implies that what is right is right in virtue of natural facts that have normative force independently of whatever motivation (or other mental stance) obtains in the situation.[50] Since I understand moral normativity and moral knowledge in terms of motivation within groups of animals, moral objectivity in this essay should not be taken to entail this stronger version of moral realism.

It is conceivable that Alice has this stronger version of moral realism in mind when she says that the difference between right and wrong is "beyond each specific culture", but it is equally conceivable that she did not. The position that I defend is less ontologically demanding than the stronger moral realism noted but entirely adequate, I would argue, to account for the objective nature of animal moral normativity and animal moral knowledge.[51]

What about the privileged status that Alice accords to love? Morality extends beyond immediate family and close friends, as we have noted. Alice must surely grant this truth that holds in the case of animals protecting those more vulnerable in the face of external danger, whether loved or not, but I think that she would. How then does love have a privileged status?

Her point, as I interpret it, is that moral norms that are central to love are likewise central to morality and moral knowledge. Two norms central to loving others are protecting them from harm and being loyal to them. Each of these norms, as Alice could easily concede, applies beyond loved ones but remains central to morality. Central in what way? The emotions of sympathy and loyalty that sustain them appear early in individual animal development and without them other moral norms arguably would not have evolved.[52] Except for them, would anything else matter? It is hard to imagine it would. Why would we want to be fair to others or respect their liberty if we feel no loyalty to them and don't care whether they are harmed? That, I submit, is

Alice's point about the status of love. Though the difference between moral right and wrong is in an important sense objective and universal, extending beyond love, without the norms of care and loyalty—the norms that pertain to love—none of the other norms that we have would have evolved or inspired separate normative commitment in the absence of love. Nor would the basic norms of care and loyalty and the emotions that sustain them have evolved to reach beyond those we love.

Conclusion

We began with familiar sentiments about the objectivity and universality of moral knowledge coupled with the sentiment that morality is anchored in our capacity for loving others. Many of us would endorse all three sentiments. We then reviewed reasons why such sentiments appear to be in direct conflict with thinking, as many of us do, that morality evolved through Darwinian natural selection. The main burden of this essay is to show how the puzzles dissolve when we shift from a human-centered conception of morality to Darwin's animal-centered understanding of moral knowledge. From that perspective, moral know-how evolved in other mammals through natural selection long before humans existed. Though moral knowledge can be more intellectually sophisticated in humans, moral knowledge for both is anchored in animal sociality, not abstract principles. Moreover, while moral norms evolve and at times regress morally, their function to sustain interdependent living explains their existence. In closing, I must caution, however, that nothing in this picture guarantees that moral norms and moral knowledge will not become extinct or that mammalian life will not disappear from the earth because humans collectively have not acted on the moral knowledge that they have. The norm-independent conditions that make possible moral knowledge do not guarantee that moral knowledge will guide our choices as we try to cope with the existential threats before us.[53]

[1] Rooney, Sally, *Beautiful World, Where Are You* (Farrar, Straus, and Giroux, 2021), Chapter 29.

[2] To say that the difference between right and wrong "transcends" culture (etc.) could imply that truths about right and wrong are intrinsically normative moral truths whose existence is independent of the evolution of moral norms (a strong form of moral realism). However, it could mean instead that the moral difference between right and wrong is fully explained by non-normative truths about how moral norms evolve where these truths are not themselves moral norms. I will take advantage of this ambiguity in the sections to follow to resolve the apparent puzzle. When the puzzle is interpreted the first way, the puzzle reflects a genuine contradiction (but see note 17). However, when transcendence is interpreted the second way, the puzzle is resolvable since the norm-*dependent* evolution of moral norms is compatible with norm-*independent* non-normative truths about the evolution of moral norms that explain their evolutionary function.

[3] For elaborations of this apparent incompatibility, see: Michael Ruse, *Taking Darwin Seriously* (Basil Blackwell 1989), Richard Joyce, *The Evolution of Morality*, (MIT Press, 2006), Sharon Street, "A Darwinian Dilemma for Realist Theories of Value" *Philosophical Studies* 127 (1) (2006): 109-166, and Guy Kahane, "Evolutionary Debunking Arguments" *Nous* 45 (1) (2011): 103-125.

[4] Jonathan Haidt and Craig Joseph, "The Moral Mind: How Five Sets of Innate Intuitions Guide the Development of Many Cultural-Specific Virtues and Perhaps Even Modules," in *The Innate Mind Volume 3: Foundations and Future*, ed. P. Carruthers, S. Laurence, and S. Stich (Oxford University Press, 2008), 367-91; Jonathan Haidt, *The Righteous Mind: Why Good People Are Divided by Politics and Religion* (Vintage Books 2012); and Victor Kumar and Richmond Campbell, *A Better Ape: The Evolution of the Moral Mind and How it Made us Human* (Oxford University Press, 2022): 83-103.

[5] David B. Wong, *Natural Moralities* (Oxford University Press, 2006).

[6] As in the case of the first puzzle, there is a hidden ambiguity in what seems to be implied in the statement of the puzzle. The difference between right and wrong may reside not in universal moral norms, as the statement of the puzzle seems to imply, but in evolutionary truths that are not themselves universal moral norms. The difference between right and wrong understood in the second way is, I shall argue, entirely compatible with moral pluralism.

[7] Unlike the supposed transcendent and universal nature of morality, the centrality of love is not normally a target for debunking based on the evolution of morality by natural selection. That is because love is not normally thought to be at the core of all morality. Alice's view of morality is different in this respect but no less at odds with the way that morality is thought to evolve by natural selection. It is interesting that Alice's challenging perspective is vindicated by a deeper understanding of the evolutionary origin of morality, as I will try to demonstrate.

[8] Michael Tomasello, *A Natural History of Human Morality* (Harvard University Press, 2016).

[9] Charles Darwin, *The Descent of Man, and Selection in Relation to Sex* (Princeton University Press, 1871/1981), 105; see also 35, 106.

[10] Darwin did not, however, believe that animals lacking mental powers equal to those in humans would have a moral conscience or be moral beings in the sense that humans are (Darwin, *The Descent of Man*, Chapter 4). His animal-centered approach to the evolution of morality is, in this respect, less thoroughgoing than the one that I defend in this essay.

[11] For example, Russ Shaffer-Landau, "Evolutionary Debunking, Moral Realism, and Moral Knowledge" *Ethics and Social Philosophy* 7 (1) (2012): 1-37; Derek Parfit, *On What Matters* (Oxford University Press, 2017); Micah E. Lott, "Must Realists be Skeptics? An Aristotelean Reply to a Darwinian Dilemma" *Philosophical Studies* 175 (2018): 71-96.

[12] Mark Rowlands provides a useful analysis of this traditional opposition in Mark Rowlands, *Can Animals Be Moral?* (Oxford University Press, 2012), Chapters 1-4.

[13] Mark Rowlands, *Can Animals Be Moral?* See also Frans de Waal, *Good Natured: The Origin of Right and Wrong in Humans and Other* Animals (Harvard University Press, 1996) and Marc Bekoff and Jessica Pierce, *Wild Justice, The Moral Lives of* Animals (University of Chicago Press, 2009). More recent work is cited in note 15. These studies like mine focus on the capacity of animals to act on moral norms.

[14] A standard theory about behavioral modernity, first articulated by Richard Klein in "Anatomy, Behavior, and Modern Human Origins" *Journal of World Prehistory* 9 (2) (1995): 167-98, is that between 50,000 and 100,000 years ago a novel neural mutation gave Homo Sapiens the capacity for "symbolic thought" and thus made them fundamentally different in thinking and feeling from all other animals. For the reasons given by Kumar and Campbell, *A Better Ape,* 145-147, this assumed cognitive divide between animal and human evolution is untenable. Significantly, the Klein story fails to make genetic sense given what we now know about when Sapiens began to emerge in sub-Saharan Africa about 300,000 years ago.

[15] Sarah Vincent, Rebecca King, and Kristin Andrews, "Normative Practices of Other Animals" in ed. Aaron Zimmerman, Karen Jones, and Mark Timmons, *The Routledge Handbook of Moral Epistemology* (Routledge, 2018), 5783. For a recent defense of the ability of animals to exercise more than instrumental rationality, see Giacomo Melis and Susana Monso, "Are Humans the Only Rational Animals?" *The Philosophical Quarterly* (2023): 1-21.

[16] As noted earlier, some will object that this approach to moral normativity runs the risk of ignoring the intellectual sophistication needed for morality to evolve. For two examples: Tomasello argues that to be moral, cooperators must be able to identify the mental states of each other and know that they are moral—a task beyond other animals (Tomasello, 2016); Korsgaard argues that to be moral animals must have the capacity of "normative self-governance" that allows them to question the propriety of their own motives (Christine Korsgaard, "Reflections on the Evidence of Morality," *The Amherst Lecture in Philosophy* 5 (2010), 1-29). These are valid points if morality is viewed from a human-centered perspective. In the present essay, inspired by Darwin, I aim for an evolutionary perspective that can enable us to understand why moral animal behavior evolved in the first place. By analogy, some eyes are so primitive that they may seem not to be eyes at all, but if they enable an animal to see well enough to survive, they function as eyes and thus earn the label.

[17] This appearance is in tension with important work by John Collier and Michael Stingl, *Evolutionary Moral Realism* (Routledge, 2020). They interpret animal morality so that moral facts are stance independent. But, if this realist view is right, my vindication of Alice's sentiments will not be undermined. Rather this moral realist interpretation of her sentiments will be defensible that is an alternative to my less demanding, moderately realist, understanding of her sentiments. In this event, the puzzles for Alice could be resolved either way. I do not need, therefore, to evaluate the animal-centered, evolutionary moral realism of Collier and Stingl to defend a much less demanding resolution of the puzzles. It is enough to note that an alternative approach may be also viable.

[18] Kristin Andrews, "Naïve Normativity: The Social Foundation of Moral Cognition" *Journal of the American Philosophical Association* 6 (1) (2020): 36-56.

[19] Many non-human animals form protective circles, such as dolphins, musk ox, bison, and elephants. Bison have been seen to form two concentric circles, with the females around their young and the males around the females. Norms for forming protective circles around the vulnerable have two key moral features: motivation to conform to the norm includes moral feelings, such as sympathy and loyalty, and conforming promotes group survival through interdependent living—the function of moral norms. The evolutionary function of moral norms and their motivation are elaborated in this section.

[20] For discussion of the evidence, see Andrews, 2020, and Vincent *et al* (2018). The latter provides tables listing numerous cases of naïve normativity in chimps and cetaceans. Examples include norms of obedience, reciprocity, care, responsibility, and solidarity.

[21] Andrews, 2020.

[22] Vincent *et al* (2018).

[23] Philip Kitcher, "Function and Design" in ed. David Hull and Michael Ruse, *The Philosophy of Biology* (Oxford University Press, 1998), 158-179. There is a complex literature around the concept of function applied to evolution, beginning with Larry Wright, "Functions" *The Philosophical Review* (1973) 82: 139-168. Since the ensuing complications do not undermine my appeal to "selected effect" to understand the function of moral norms, I ignore them. I should note that Jeffrey Wisdom appeals to moral function to defend moral realism in "Proper Function in Moral Realism" *European Journal of Philosophy* (2017) 25 (4): 1660-1674. In contrast to the present paper, however, Wisdom's approach is human centered.

[24] For an extended defense, see Kumar and Campbell, 4-6, 11, 38, 55, 58, 80, 83, 89, 94, 96-99, 106, 116, 121, and 123-124.

[25] Kumar and Campbell, 127-171.

[26] Kumar and Campbell, 196-199.

[27] G.E.M. Anscombe, *Intention* (Cambridge, MA: Harvard University Press, 1957).

[28] The example is from Ruth Millikan, "Pushmi-Pullyu Representations" *Philosophical Perspectives* 9 (1995): 185-200.

[29] For vivid examples, see Alan Fiske and Tage Shakti Rai, *Virtuous Violence: Hurting and Killing to Create, Sustain, End, and Honor Social Relationships* (Cambridge University Press, 2014); Jonathan Glover, *Humanity: A Moral History of the Twentieth Century* (Yale University Press, 2000).

[30] Richmond Campbell and Jennifer Woodrow, "Why Moore's Open Question Is Open: The Evolution of Moral Supervenience" *The Journal of Value Inquiry* 37 (3) (2003): 353-372.

[31] See note 19 for some examples of animals that form protective circles.

[32] An understanding of moral judgment that allows this possibility is in Richmond Campbell, "What Is Moral Judgment?" *The Journal of Philosophy* 104 (7) (2007): 321-349; see 333-335 for discussion of the example.

[33] Eve Roberts pointed out in personal correspondence that this way of conceiving moral knowledge has a distinctly Aristotelian favor. I agree, though Aristotle has a more demanding conception of moral agency, as noted earlier.

[34] For elaboration and defense of moral consistency reasoning as it is conceived in this essay, see: Richmond Campbell and Victor Kumar, "Moral Reasoning on the Ground" *Ethics* 122 (2) (2012): 273-312; Richmond Campbell and Victor Kumar, "Pragmatic Naturalism and Moral Objectivity" *Analysis* 73 (3) (2013): 446-455; Richmond Campbell, "Learning from Moral Inconsistency" *Cognition* 167 (2017): 46-52; and Kumar and Campbell (2022), 109-121.

[35] For examples of interdependent living within groups of elephants, wolves, and killer whales that appears guided by what I call moral knowledge, see Carl Safina, *Beyond Words: What Animals Think and Feel* (Henry Holt, 2015).

[36] The claim that human moral knowledge is fundamentally different in kind is made in Paul Thagaard, "Darwin and the Golden Rule" *Biology and Philosophy* 37 (2022): 58-76.

[37] Safina, 18.

[38] Safina notes that primates know third party relationships, and so do wolves, hyenas, dolphins, and birds of the crow family. Might moral norms about proper treatment of family be reinforced for these animals too based on third party understanding of family relations? It would not be surprising if they were, though my conclusions on moral normativity among animals do not depend on this possibility.

[39] It might be objected that the elephants follow the matriarch from habit rather than knowing morally what to do. While knowing morally how to act in a situation can become mostly a matter of habit, it is never just a matter of habit. When the circumstances change in a morally relevant way, what is morally required can be different. In the previous example, Tecla stopped her habitual routine, however moral, to object to the other mother's neglectful behavior. The very young will normally act merely from habit in following the matriarch until they don't and learn from the reaction of others the cost of deviating from the norm. Afterward, their norm following ceases to be merely habitual.

[40] For elaboration of this response to the alleged counterexamples of slavery, sexism, and speciesism, see Kumar and Campbell (2022), 175-251; to pursue them here would take us too far afield.

[41] I owe this objection to Duncan MacIntosh.

[42] Street, "A Darwinian Dilemma".

[43] This consideration blocks the idea that the Darwinian origin of moral norms entails Social Darwinism, the view that we ought morally to promote the survival of the fittest. That directive is too general to function as a moral norm. Moreover, it would be morally inconsistent with core moral norms that do function morally and, ironically, do so precisely because of natural selection.

[44] Robert Brandon presents these challenges in his, "Adaptation and Evolutionary Theory" *Studies in the History of Philosophy of Science* 9 (1978): 181-206; reprinted in Robert Brandon, *Concepts and Methods in Evolutionary Biology* (Cambridge University Press, 1996), 3-29.

[45] See, for example, Richmond Campbell and Jason Robert, "The Structure of Evolution by Natural Selection" *Biology and Philosophy* 20 (2005): 673-696.

[46] See note 4.

[47] See Campbell and Kumar, "Moral Reasoning on the Ground" Section V, and Kumar and Campbell (2022), Chapter 5, for the importance of moral consistency for the successful functioning of moral norms in the human case. The argument in the case of other animals would be parallel.

[48] See note 24.

[49] For an understanding of how realist moral naturalism is a form of moral realism, see: Sayre-McCord, Geoffrey, "Moral Realism", *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/win2023/entries/moral-realism/>. A notable recent example of realist moral naturalism is William Rottschaefer, "Affording Affordance Moral Realism" *Biological Theory* 16 (2021): 30-48.

[50] See, for example, the characterization of moral realism in Sharon Street, "A Darwinian Dilemma". Does the moral realism in Rottschaefer, "Affording Affordance Moral Realism" constitute a naturalism that exemplifies this stronger form of moral realism? Does the naturalism in Collier and Stingl, *Evolutionary Moral Realism*? Fortunately, the realist moral naturalism defended in this essay does not require answers to these difficult questions.

[51] For an in depth and thorough discussion of this issue, see David Copp, *Ethical Naturalism and the Problem of Moral Normativity* (Oxford University Press forthcoming).

[52] For details see Kumar and Campbell (2022), Chapters 1-2.

[53] For an elaboration of this theme see W. Ford Doolittle, "Could this Pandemic Usher in Evolution's Next Major Transition?" *Current Biology* 30 (15) (2020): R846-848.

Richmond Campbell
Dalhousie University
richmond.campbell@dal.ca