

4-14-2017

The Epistemic Status of Moral Conceptual Truths

Kara D. Boschert

University of Missouri-St. Louis, kdbq6b@mail.umsl.edu

Follow this and additional works at: <http://irl.umsl.edu/thesis>

 Part of the [Biology Commons](#), [Epistemology Commons](#), [Evolution Commons](#), and the [Other Philosophy Commons](#)

Recommended Citation

Boschert, Kara D., "The Epistemic Status of Moral Conceptual Truths" (2017). *Theses*. 293.
<http://irl.umsl.edu/thesis/293>

This Thesis is brought to you for free and open access by the Graduate Works at IRL @ UMSL. It has been accepted for inclusion in Theses by an authorized administrator of IRL @ UMSL. For more information, please contact marvinh@umsl.edu.

THE EPISTEMIC STATUS OF MORAL CONCEPTUAL TRUTHS

Kara D. Boschert
University of Missouri-St. Louis
B.A. Philosophy, Truman State University 2014

A Thesis Submitted to The Graduate School at the University of Missouri-St. Louis in
partial fulfillment of the requirements for the degree
Masters of Arts in Philosophy

May
2017

Advisory Committee

Billy Dunaway, Ph.D
Chairperson

Lauren Olin, Ph.D
Co-chair

Eric Wiland, Ph.D
Co-chair

Copyright Kara D. Boschert

Introduction

Evolutionary debunking arguments assume that morality could, conceptually speaking, be about anything. One response to this contention is that there are some moral conceptual truths which counter assertions that we could be deeply in error about basic moral truths. I argue that moral conceptual truths do not meet the minimal epistemic requirements needed to show that debunkers are wrong about the conceptual flexibility of morality. Specifically, moral conceptual truths do not meet basic safety conditions which preclude luckily true beliefs counting as knowledge.

Here is an outline of how this paper will progress. First I detail evolutionary debunking arguments and explain how moral conceptual truths could be considered a promising response to debunkers. Next, I clarify epistemic safety as a basic consideration in any theory of knowledge. After that, I outline a plausible theory of the evolution of human morality, providing some of the biological mechanisms from which a human moral sense likely developed. In light of the evolutionary development of human morality, I will provide scenarios that show our moral concepts could have easily differed from the ones we actually have. I conclude that moral conceptual truths do not meet safety conditions and thus do not constitute a successful response to the debunker's premise that morality could be about anything.

Debunking Arguments and The Moral Fixed Points

Evolutionary Debunking Arguments

Evolutionary debunking arguments aim to show that we are in moral error because our moral beliefs are products of evolution. Evolutionary debunking arguments can be characterized in five steps:

1. The realist says there are attitude-independent moral truths.
2. Evolutionary forces have influenced our beliefs.
3. Evolution selected for beliefs that increased reproductive success rather than the tracking of independent moral truths.
4. The moral truths and the beliefs which promote fitness come apart.
5. So, our moral beliefs are probably mistaken. (Vavova 5)

A substantive assumption that debunkers make in this argument is that morality could, conceptually, be about anything. Evolution is going to select for survival-promoting beliefs, but the debunkers' hunch is that morality could actually be about something else (Vavova 112). Even if morality *were* about doing fitness-decreasing things, like killing our children, evolution would still select for the belief that we should not kill our offspring, and we would be in grave moral error. Vavova puts the debunker's contention like this: "It is conceptually possible that morality is about throwing ourselves off cliffs and causing ourselves pain. If it had been, evolution would have still inclined us to think it was about survival and pain avoidance" (Vavova 112). The conclusion the debunker draws is that we cannot trust seemingly obvious moral beliefs.

Sharon Street, known for her often-cited version of the debunking argument, “A Darwinian Dilemma for Realist Theories of Value,” also discusses the open conceptual nature of morality. Street is quite strong in her contention that morality as well as our basic moral concepts such as pain, are conceptually flexible. Street holds that appeals to the undeniable badness of pain are not good responses to debunking arguments because pain, in theory, could be a *good* thing. Street asserts the following concerning the conceptual possibility of pain not being bad: “...it is perfectly possible, as a conceptual matter, that instead of disliking [the] pain the way we all happen to do, we could naturally enjoy it and seek it out...” (Street 40). Street asserts pain, as a matter of conceptual possibility, could be a good thing. Even if it were good, we would still falsely believe pain is bad because such a belief is fitness-promoting (Street 43).

Moral Conceptual Truths as a Response to Debunking Arguments

Before I talk about *moral* conceptual truths, I should explain my favored understanding of conceptual truths. Some stock examples of conceptual truths are propositions like, “all bachelors are unmarried males,” and “a vixen is a female fox.” Negating propositions like these is considered conceptually impossible (Williamson 1). The conceptual impossibility of denying these propositions has to do with the conceptual connections concepts like “vixen” and “bachelor” have. For instance, bachelor has special conceptual connections to the concepts “unmarried” and “male.” Another way of putting this idea is that “male” and “unmarried” are constituent concepts that are built in to the composite concept “bachelor” (Williamson 1). Conceptual truths have unique

epistemic features. A common view is that understanding a conceptual truth is enough to know that it is true (Williamson 3). For example, if I understand the proposition, “all bachelors are unmarried males, I also know that proposition is true.

Vavova talks about the relevance of moral conceptual truths as a response for realists against debunkers, but does not develop a response herself (Vavova 112). Other work has been done, though, to show that moral conceptual truths offer realists an auspicious comeback against debunkers. Moral conceptual truths allow realists to assert that, contrary to the debunker’s sentiment, morality cannot be about *anything*. Proponents of moral conceptual truths hold there are special truths which establish conceptual limits to morality. According to proponents of moral conceptual truths, some things, by definition, cannot count as moral.

Terence Cuneo and Russ Shafer-Landau have curated a list of plausible moral conceptual truths called “The Moral Fixed Points” and argue that the fixed points establish conceptual limits to morality. Among the candidates for moral fixed points that Cuneo and Shafer-Landau offer are propositions such as, “It is pro tanto wrong to engage in the recreational slaughter of a fellow person,” “It is pro tanto wrong to break a promise on which another is relying on simply for convenience sake,” and “It is pro tanto wrong to torture others just because they have inconvenienced you” (Cuneo & Shafer-Landau 405). Because the moral fixed points are conceptual truths, anyone who rejects them is not capable of moral reasoning (Cuneo & Shafer-Landau 401).

While it is possible you may have qualms with a few of Cuneo and Shafer-Landau’s candidates for moral conceptual truths, we should acknowledge that Cuneo and

Shafer-Landau only suggest a list of *candidates* for moral conceptual truths. We need not accept their list of moral fixed points wholesale to be on board with their general project. Cuneo and Shafer-Landau's aim is to show that there are conceptual truths that are at least similar the ones they offer, although it's possible the official list could look a little different. The relevant point is that there are *some* conceptual limits to morality, so it is not 'anything goes' when it comes to what could, conceptually, be considered moral. For example, if you believe that wanton killing is morally good, you lack the concept of morality.

We can see how if there are some moral conceptual truths like Cuneo and Shafer-Landau's moral fixed points, the debunker's contention that morality could be about anything is defeated. The debunker maintains that it is conceptually possible for morality to be about infanticide or causing pain, so it is possible that we are in deep moral error. If the *pro tanto* wrongness of wanton killing and torture are conceptual truths, then the debunker's argument lacks force. It is not possible that we could be in the acute moral error the debunker has in mind if our most central and foundational moral beliefs are true as a matter of conceptual necessity.

As mentioned earlier, Sharon Street maintains that pain, conceptually speaking, could be good. This is quite a bold claim as it entails that we are possibly mistaken about the badness of pain. This, simply put, sounds crazy! Moral conceptual truths help to explain why Street's claims about the conceptual flexibility of pain seem obviously wrong. A proponent of moral conceptual truths can respond to Street by saying that pain, by definition, is bad; one lacks the right concept badness if one denies that pain is bad.

The moral fixed points offer realists a promising response to a premise in debunking arguments¹. In addition, one could also argue they provide a solid bedrock for our moral knowledge. The moral fixed points are not a random list of moral truths; the fixed points are truths that are central and foundational to our moral knowledge. Our most cherished moral beliefs, such as the wrongness of killing, stealing, cheating, etc, as well as the beliefs they support, are safe from evolutionary challenges. Further, if foundational moral beliefs like “killing is wrong” or “stealing is wrong” are not subject to evolutionary debunking worries, then the possibility of error is mitigated to our more peripheral moral beliefs. The overall upshot is that our knowledge of central moral judgements such as the wrongness of cheating or killing is secured. It is only our more peripheral, contextual moral beliefs (i.e. It is wrong to take a screenshot of a personal text message and share it with another party) that remain subject to the debunker’s worries.

Epistemic Luck and Safety

Epistemic luck and the related notion of epistemic risk remain factors that undermine knowledge no matter what one’s favored understanding of justification. A familiar instance of luck is Russell’s example of a man looking at a broken clock at just the right moment. The moment when the man decided to look at the clock just happened to be the moment when the clock’s hands were pointing to the correct time. Had the man glanced at the clock a minute later, the man would have formed a false belief about what

¹ The premise that morality could be about anything supports the debunker’s contention that it is possible that we are in deep moral error.

time it is. In the broken clock scenario, the man lacks knowledge of the time in spite of his true belief because his belief was lucky. (Russell 170).

Epistemic luck also seems to be relevant to a diagnosis of so-called Gettier problems. Edmund Gettier famously showed how one's belief could be true and justified yet still not qualify as knowledge. One of Gettier's examples is a man named Jones who applied for a job and has strong evidence that he will get the job. In addition, Jones knows he has ten coins in his pocket. Based on robust evidence that he will get the job, Jones believes the conjunct "Jones will get the job and Jones has ten coins in his pocket." Since that conjunct entails that the man who will get the job has ten coins in his pocket, Jones also believes "the man who will get the job has ten coins in his pocket." Unfortunately for Jones, a man named Smith got the job. Also unbeknownst to Jones, Smith happens to have ten coins in his pocket. (Gettier 122) Jones's belief "the man who has ten coins in his pocket" is both justified and true, but only by luck. We do not count Jones's belief as knowledge because he only got things right by chance.

Luckily true beliefs are insidious because they seem to undermine knowledge in all cases, so luck is relevant regardless of the epistemology one accepts. Since luck is related to core epistemological notions like knowledge and justification, luck is everyone's problem. For example, externalists might think about justification in terms of reliable processes, while internalists focus on whether there is a defeater for a belief, but both internalist and externalist versions of knowledge and justification are compromised by the presence of luck or chance. If a belief is only luckily true, it does not matter if the belief was produced by a reliable process or that there are no defeaters for it. Since the

belief could have very easily been false, something epistemically has gone awry.

Everyone can be on board with the idea that luck is a problem for knowledge. In light of the ubiquity of luck, the following is a consideration that any theory will want to include.

One way of refining our understanding of epistemic luck is through the notion of epistemic safety. Safety is the idea that a belief that could have easily gone wrong cannot count as knowledge. Another way of putting that is if you know p , your belief about p could not have easily been false. David Manley uses the example of hearing a bird call to illustrate an unsafe belief (Manley 403). If he thinks he hears a lark call outside and forms the belief, “that’s a lark,” his belief is not safe if there are lark-imitators nearby. After all, even if the call he heard came from a real lark, it could have easily been a false lark call. The belief “that is a lark” could have easily been false.

How do we determine whether a belief could have easily gone wrong? One way of regimenting our notions of luck and safety is to talk in terms of nearby worlds. If we adopt nearby world talk, being lucky means there is a nearby world where you falsely believe. The closer the world where you mess up, the less safe you are. Conversely, the more distant the worlds are where you get thing wrong, the safer you are. (Manley 406) Setting the issue of which worlds count as ‘closer’ aside for the moment, we can consider the general conditions under which belief in a moral conceptual truth is safe. In order to safely believe in a moral conceptual truth (so, in order to know it assuming all other epistemic standards are met), the belief could not have easily gone wrong in a close-by world.

Manley points out how it is not just risk of error about a belief p , but also propositions related to p , that can undermine knowledge. If falsely believing p in a nearby world makes one's belief unsafe, then so does falsely believing propositions that are similar to p (Manley 406). Manley calls this amendment *revised safety*. Given our revised conceptions of safety, similar untrue thoughts in nearby worlds count against knowledge. Very different untrue thoughts in closer worlds do not count against knowledge and neither do very similar thoughts in distant worlds. (Manley 407)

Here is how revised safety could fail to be met based on falsely believing a related proposition. Note that this is only one example of a non-safe belief based on revised safety, and there are plenty of other ways that a belief could be unsafe. My aim is to clearly illustrate the conditions of revised safety through an example of how its conditions could fail to be met. Here is the example: My belief is not safe if I chose to use an adjective to describe something, but I could have easily picked a similar adjective, with slightly different conceptual entailments, which would have led to a false belief. The following is a specific example of how that could happen. Be aware that I assume that beast and animal have different conceptual entailments. I am on a hike in the woods and I see a dog chewing on a bone and I form the belief, "that animal is hungry." My belief is true, however I could have easily² instead formed the belief, "that beast is hungry." Unbeknownst to me, the dog is a lost pet and therefore would not meet the conceptual entailments of beast-hood (i.e. being disposed to untamed, aggressive behavior). So, by

² Read: In a nearby world I would have formed the belief "that beast is hungry."

using the concept beat instead of animal, I would have either formed a false belief or have been conceptually committed to having false beliefs.

Morality's Evolutionary Genesis

There are many plausible evolutionary accounts of the origins of human morality. I will be focusing on an evolutionary account of morality which understands the human moral sense as an adaptation to promote social cooperation. According to the account I favor, morality plays a centrally social role. Non-cooperative, anti-social behaviors like cheating and stealing were likely the first behaviors that generated moral prohibitions. Because moral prohibitions against cheating and stealing mark the origin of human morality, judgments about the wrongness of such behaviors can be considered foundational moral judgments, at least in the evolutionary sense.

Biological Origins of Caring and Helping

When telling an evolutionary story about human morality, one must note the conservative character of the biological process of natural selection. Co-opting old structures for new purposes is more economical than generating entirely new ones. In the case of human morality, all of the requisite mechanisms and structures were in place for a moral sense to develop. The hyper-social nature of human beings created the need for a new social regulation process, and a moral sense became adaptive. More ancient mechanisms were re-appropriated for the purposes of facilitating human social engagements, marking the origin of human morality.

Before I elaborate further about which social problems morality evolved to solve, I will describe some of the previously existing pro-social mechanisms which provided the groundwork for human morality in the form of promoting cooperative and caring behaviors. Kin selection is among one of the most ancient mechanisms believed to provide some of the basic structures for the human moral sense because it looks to be the genesis of helping and caring behaviors. Kin selection is the natural selection of creatures which tended to help their offspring and relatives (Joyce 32). Caring and helping behaviors toward an organism's kin is undoubtedly adaptive behavior; genes which promote ensuring the survival of other creatures with share those genes are more likely to prosper (Joyce 32).

Mutualism, another mechanism which promoted the development of morality, involves the cooperation of creatures in a potentially mutually-beneficial situation. For instance, a lion pride that works together to kill a giraffe will all benefit from cooperative behavior since none of them could bring down a giraffe alone (Joyce 22). Another instance of helping behavior promoted by evolution is direct reciprocity, also known as reciprocal altruism. Direct reciprocity promotes a short-term sacrifice in exchange for a long-term gain. A classic example of direct reciprocity are coral reef "cleaning stations" where shrimp clean the inside of a larger fish's mouth from parasites and pollutants. Both the larger fish and the shrimp make a temporary sacrifice in this scenario. The larger fish sacrifices an easy lunch because it could easily swallow a mouthful of the shrimp, while the shrimp sacrifice their safety. In spite of the temporary decrease in each other's fitness, both creatures benefit from these exchanges in the long run. The larger fish remains free

from parasites and pollutants which could shorten its lifespan, and the shrimp get a regular source of food. (Joyce 24)

Human Morality

Although the caring and cooperative behaviors described above can be considered the evolutionary foundations for human morality, they are not examples of morality. Morality, among other things, is known for its categorical and authoritative character, features that I will explain later. Helping behavior alone does not constitute a moral sense. According to the empirical accounts I favor, natural selection co-opted preexisting mechanisms like kin selection and direct reciprocity and appropriated them to solve unique problems that arose from human sociality.

Morality evolved because it helped facilitate social interactions specific to humans. Richard Joyce asserts that morality was particularly crucial for the maintenance of reciprocal relations among humans (Joyce 140). Similarly, Allan Gibbard holds that a moral system evolved to facilitate social coordination (Gibbard 61). Morality was advantageous to humans because it provided a mechanism that reliably prohibited anti-social behaviors like cheating and stealing. Here is a specific example, inspired by Gibbard and Joyce, of the type of social scenario that the inclusion of morality facilitated. According to the accounts I have in mind, morality evolved to facilitate exchanges between humans in mutually beneficial scenarios. Cooperative behavior on a hunt is necessary for the survival of each member of the group. One man cannot hunt and bring down a buffalo easily, so hunting in a group is cost-efficient for everyone. In order for all

participants to mutually gain from the mutual endeavor, imperatives that enforce fairness and cooperative behavior are necessary (Gibbard 67).

Morality does not just amount to a desire to maintain pro-social behaviors. Morality, at its most sparse, provides universal prohibitions which is what makes it so effective at promoting the right behaviors. Morality was a useful tool for coordinating human social interactions because moral judgments purport to be more than mere statements of negative or positive affect. Moral judgments are both authoritative and categorical. The authority of moral judgments means that moral considerations always win out against other considerations, making morality an effective way at promoting pro-social behavior. In addition, the categoricity of morality makes morality's commands universal. This categorical aspect of morality ensures the individual that morality applies to all humans. The categoricity of morality obligates everyone to reciprocate helping behavior. Knowing others will eventually reciprocate motivates one further to help even when help will not be returned right away.

These authoritative and categorical features of morality are what make morality an evolutionarily useful tool. Morality is uniquely both an action-guiding mechanism and a signaling tool. Morality promotes evolutionarily advantageous pro-social behaviors, but it also allows one to signal to others that one will cooperate. The upshot is that humans will often be motivated to help and cooperate expecting nothing in exchange. This unrewarded pro-social behavior assures other members of the group that one is good to have around. Moral behavior, even if not immediately rewarded, is an investment that will be likely rewarded in the future. Others are more likely to partner with someone who

has proven they are helpful and trustworthy. In addition, it is good for everyone that helpful members of the group be kept around; helping others with nothing in exchange is a way to earn safety in the long-run.

Moral Emotions and Moral Concepts

I have provided a general sketch of an evolutionary account of human morality. The story breaks down further concerning moral concept acquisition. Evolution endowed us with a set of emotional tendencies which inform our normative reasoning. Moral emotions are key to explaining the mechanism of the punishment-reward system that is built into human morality that promotes pro-social behaviors. Moral emotions allow humans to signal to one another that we have morality. If I know that you have morality, I know I can trust you not to cheat, for instance, because you will feel guilty if you do. There is a penalty (in the form of an unpleasant emotion) that accompanies anti-social behaviors. (Gibbard 139) You might reasonably wonder why the story does not stop there. If we have moral emotions, isn't that all we need to generate evolutionarily advantageous pro-social behaviors? Although I have granted the existence of moral concepts from the outset of this paper, I think there is a plausible argument for the necessity of moral concepts available.

A representationalist account of beliefs and desires can explain why we have moral concepts in addition to moral emotions. For the sake of brevity and credibility, I am adopting Elliott Sober and David Wilson's account of representationalism with respect to beliefs and desires. Moral emotions, like other sorts of emotions, cause us to

have certain motives, beliefs or desires. States like beliefs and desires are propositional attitudes, meaning they have propositional content. For instance, if I have a desire for there to be beer in my cup, I have a certain relation with the proposition “there is beer in Kara’s cup.” Specifically, I have a representational relation with the proposition contained in my desire. Believing or desiring requires the formation of a mental representation of propositions. (Sober & Wilson 208, 209)

In order to be capable of representing a proposition, an organism must have the constituent concepts that the proposition contains. In my case, I have to have the concepts ‘beer’ and ‘cup’ in order to have the desire that there be beer in my cup. This view does entail that possession of a natural language is necessary for forming beliefs and desires. Because beliefs and desires represent propositions, they contain concepts. Without possession of a language, a creature cannot have concepts³. (Sober & Wilson 209)

Given the story I just told about the relation of moral judgments with moral concepts, we can see how evolutionary forces indirectly influence our moral concept formation. Moral emotions help motivate cooperative behavior, but our moral emotions have also endowed us with certain moral concepts. Our moral concepts are contained in the proposition that our moral desires and beliefs represent. If our moral emotions have been directly selected for, then our moral concepts have been indirectly selected for insofar as they are dependent on our moral emotions producing particular beliefs and desires.

³ The bullet a proponent of this view has to bite is that animals (with the exception of animals that possess language), strictly speaking, do not have beliefs, desires, or concepts. I am willing to accept this consequence because we can still say that an animal is having representations which may be similar to beliefs or desires in humans, but animals do not, strictly speaking, have beliefs and desires.

Are Moral Concepts Safe?

Given the empirical accounts of moral concept acquisition that I favor, I am able to tell an alternative story of how we could have obtained moral concepts with conceptual entailments that are slightly different than the moral concepts that we have in this world. That we would have acquired different moral concepts if the conditions under which our concepts developed were different is not a new idea. In “An Enquiry into the Source of Morals,” Hume asks us to imagine a world of abundant resources where no hard work was required and everything humans needed was so abundant that everyone could have as much as they wanted. In such a world, Hume tells us, the concept of justice would be useless, so we would not develop it (Hume 8). The world that Hume asks us to imagine is obviously very different than the actual world and would not count as a nearby world. However, the point to take away from Hume is that our moral concepts are not static. Rather, our moral concepts are products of independent factors like our biological development and environmental conditions.

In the following subsections, I provide several examples of how humans could have adopted different moral concepts in the hopes that providing several examples will better illustrate the contours of the epistemic worry I am interested in. Before I begin, note that it is important to keep in mind that just because one of the scenarios I outline does not seem as likely to occur as what *actually* happened, that does not by itself show that the alternative scenario *couldn't* happen. In addition, be aware that even if the scenarios below are not fully empirically sensitive, all I need to accomplish is indicate the

formal features of the epistemic concern at hand. Doing so does not require an absolutely thorough and empirically accurate account of the evolution of human morality.

Scenario #1: Gradiential cheating

As I discussed earlier, the wrongness of cheating is foundational to human morality in the sense that it was likely one of the first moral truths acquired. Given the basic status of the wrongness of cheating, this putative moral conceptual truths serves as a fruitful example when considering nearby-world scenarios. The scenario we are now considering asks us to imagine a world where humans have a more nuanced view of cheating than we do in this world. Overall, humans in such a world are not concerned with cheating broadly speaking like we are. In this world, humans conceptualize many types of cheating such that all forms of cheating cannot be evaluated the same. Instead of conceptualizing all forms of cheating as wrong, humans in this nearby world are likely to evaluate each form of cheating independently.

Here's a specific example of how cheating could be conceptualized in this way: cheating among relatives, which would violate kin selection, is seen as wrong. Cheating in the form of taking what one did not earn is also conceptualized as wrong. This form of cheating goes against mutualism, which involves members of the same species working together for a mutual benefit. An example would be a man who refused to participate in a hunt but still reaps the benefits of the hunt later. Finally, however, cheating in the form of violating direct reciprocity, is not conceptualized as wrong. Recall that direct reciprocity involves incurring a short-term sacrifice for a long term gain as in the coral reef cleaning

stations. We can imagine a world where forms of cheating that conflict with direct reciprocity are not conceptualized as wrong. An example might be if a man owes someone else a favor (perhaps she let him borrow her tool), but he is not willing to incur any inconveniences at the moment so he does not let her borrow his tool. The man's actions are not seen as wrong because of the subtle variances in forms of cheating that are perceived by humans in this world; his form of 'cheating' is acceptable.

The subtle variances among the conceptualized forms of cheating results in some forms of what we consider cheating being perceived as morally neutral, while other forms are morally bad. So, when considering whether cheating is wrong, the evaluation would proceed as such: cheating 1 v cheating 2 v cheating 3... etc. Instead of evaluating all types of cheating respectively, humans would instead likely form a slightly different moral concept which does not include the wrongness of cheating as one of its conceptual entailments. In this nearby world, then, cheating is wrong would not count as a moral conceptual truth based on their concept of wrongness.

Scenario #2: Other Moral Emotions

To gain insight into the former reason, let's consider again how social factors influenced the development of moral prescriptions against cheating. Moral prescriptions against cheating and stealing developed as a way to facilitate bargaining situations among our ancestors. Moral emotions were selected for as they were a way of mediating mutually beneficial situations that required cooperation, such as hunting or trading (Gibbard 65).

When considering how we could have developed variant moral conceptual truths, we can consider the role our actual moral emotions played in our MCT acquisition. If our moral emotions influenced our moral concepts, then humans with moral emotions that function differently would likely inherit different moral concepts, with variant conceptual entailments. According to Gibbard, moral emotions like guilt and anger were selected for as they facilitated social coordination (138-139). As a result, we have moral beliefs and constituent concepts that are products of these basic moral emotions. For example, the moral emotion guilt deters cheating or stealing behavior, and eventually influenced our current concept of wrongness.

Now consider an alternative course of events: we have the same moral emotions in a nearby world, but the strength that we feel them varies compared to the actual world. We can imagine a world where we feel guilt after cheating or stealing, but the guilt we feel after cheating is less strong than what we feel in the actual world. If this were so, our moral concepts would come out slightly different as our moral concepts are molded by our moral emotions. If the moral emotion that differed in strength compared to the actual world were guilt, then the anti-social behaviors which guilt regulates, such as cheating, may be conceptualized differently in that world. The concept wrong used in such a world may be similar to ours but lack the wrongness of cheating as a conceptual entailment.

Scenario #3: Pro-Attitudes Motivate

I have explained in the earlier section on the evolution of morality that con-attitudes play a large role in regulating advantageous behavior. For instance, we are deterred from anti-social behaviors like cheating and stealing because we will feel guilty if we do not cooperate. Another possible mechanism for achieving the same effect is having primarily pro-attitudes that motivate us to perform pro-social behaviors. We can imagine a nearby possible world where instead of being deterred from stealing by the prospect of guilt, we are motivated to cooperate because doing so will make us feel happy.

Of course we are partially motivated by pro-attitudes to cooperate in the actual world, but according to the evolutionary account I favor, it is the con-attitudes that are mostly responsible for producing pro-social behaviors. Sure, we cooperate and help others because it makes us feel good, but the bigger imperative is that we avoid the intensely unpleasant feelings of guilt, regret, and sorrow that we will feel if we engage in anti-social behaviors. Imagine what our moral concepts would look like if it were the other way around.

If we were primarily motivated by pro-attitudes and very little motivation came from con-attitudes, our moral concepts would have entailments that are different than in the actual world. If everyone were motivated to cooperate primarily because they would feel intense happiness or satisfaction if they did, our negative moral concepts such as WRONG or BAD would be likely more sparse, having less conceptual entailments. For example, it may be true that stealing is wrong, but the wrongness of stealing may not be a conceptual entailment of WRONG in such a world. On the flip side, our positive moral concepts like RIGHT would likely be more robust and have more conceptual entailments.

Back to Debunking

In all of the above scenarios, humans end up with slightly different moral concepts with different conceptual entailments due to a tweak in our development of morality. For example, in the first scenario, humans in a nearby world adopted a moral concept that is similar to our concept wrong except it did not entail the wrongness of cheating. So, instead of adopting our concept WRONG, inhabitants of the nearby world adopted WRONG*, which does not have the wrongness of cheating encoded in it. This variant concept was acquired through a more nuanced view of cheating than we have in the actual world. As well as being acquired differently than our concept WRONG, WRONG* functions differently. WRONG* may have many of the same conceptual entailments as WRONG so that they can be used interchangeably in some cases without resulting in a false belief. However, WRONG* lacks the conceptual entailments that WRONG has with respect to cheating.

The point to be made in light of these nearby world scenarios is that putative moral conceptual truths are not epistemically safe because we could have easily had false beliefs about what for us are considered moral conceptual truths. In all of the above scenarios, the moral concepts that humans acquire in the nearby worlds differ in the truths they encode as conceptual necessities. As a result of their variant moral concepts, humans in the nearby worlds could end up with false beliefs. Human in a nearby world may have beliefs identical to ours on the level of a proposition⁴, but since their moral concepts are

⁴ For example, they could believe “causing pain is wrong” and so do we.

slightly different than ours, they will easily end up with false beliefs. Based on their moral concepts, they could believe that stealing is wrong, however 'stealing is wrong' is not a conceptual truth based on their variant concept WRONG*. So, their concept of wrongness allows them to deny that stealing is wrong with no conceptual difficulties.

Remember, per Manley, that is it not just the close possibility of a false belief in a nearby world that threatens the safety of a belief, but it is also the close possibility of falsely believing a proposition *related* to the belief in question. The likelihood of a false belief about a moral conceptual truth in a nearby world which is closely related to our actual belief of a moral conceptual truth is enough to render our belief in the moral conceptual truth unsafe. What the scenarios above show is that we could have easily come to false beliefs about what for us are moral conceptual truths like the wrongness of cheating. Because our moral concepts could have easily been different, our beliefs in moral conceptual truths in the actual world are only luckily true at best. Moral conceptual truths, such as the moral fixed points, do not pass the safety test because we could have easily developed slightly different moral concepts which would yield false beliefs. This result renders moral conceptual truths on their own useless to the realist as a response to the debunker's premise that morality could be about anything.

Conclusion

Let me rehash what I have done in this paper. First, I explained the evolutionary debunking argument and highlighted a premise in the argument which says morality can conceptually be about anything. Next, I explained how it can be argued moral conceptual

truths refute the debunker's claim that morality can be about anything. I clarified why epistemic safety is an important consideration for any theory of knowledge. After that, I offered an evolutionary account of human morality and acquisition of moral concepts. Based on the evolutionary account of morality, I provided three nearby-world scenarios which illustrated how our moral concepts could have easily been slightly different. Finally, I put everything together and concluded that moral conceptual truths are not epistemically safe and therefore do not offer realists a successful response to a premise in the debunking argument.

The upshot is moral conceptual truths cannot help realists with a successful response to the debunker's premise that morality could be about anything because moral conceptual truths are not epistemically secure. Not all hope is lost for realists who want to establish conceptual limits to morality, but putative moral conceptual truths alone do not look like they will provide the epistemic immunity realists need in this dialogue. If realists want to deny the debunkers' assertion that morality could conceptually be about anything, they will have to find another way besides appealing to moral conceptual truths.

Works Cited

- Cuneo, Terence, and Russ Shafer-Landau. "The Moral Fixed Points: New Directions for Moral Nonnaturalism." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 171.3 (2014): 399-443.
- Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis* 23.6 (1963): 121-123.
- Gibbard, Allan. *Wise Choices, Apt Feelings*. Cambridge: Harvard University Press, 1990.
- Hume, David. *An Enquiry Concerning the Principles of Morals*, in the version presented at www.earlymoderntexts.com.
- Joyce, Richard. *The Evolution of Morality*. Cambridge: The MIT Press, 2007.
- Manley, David. "Safety, Content, Apriority, Self-Knowledge." *The Journal of Philosophy* 104 (2007): 403-423.
- Russell, Bertrand. *Human Knowledge: Its Scope and Limits*. London: Allen & Unwin, 1948.
- Street, Sharon. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 127.1 (2006): 109-66.
- Sober, Elliott and David Sloan Wilson. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge: Harvard University Press, 1998.
- Vavova, Katia. "Evolutionary Debunking of Moral Realism." *Philosophy Compass* 10.2 (2015): 104-16.

Kara Boschert
The Epistemic Status of Moral Conceptual Truths

Williamson, Timothy. "Conceptual Truth". *Proceedings of the Aristotelian Society* 80
(2006): 1-41.