University of Colorado, Boulder
CU Scholar

Philosophy Graduate Theses & Dissertations   Philosophy

Spring 1-1-2018

A Defense of the Moral Error Theory

John Alton Christmann
University of Colorado at Boulder, john.christmann@colorado.edu

Follow this and additional works at: https://scholar.colorado.edu/phil_gradetds

Part of the Applied Ethics Commons, and the Metaphysics Commons

A Defense of the Moral Error Theory

by

John Christmann

B.A. University of Pittsburgh, 2016

A Thesis Submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirement for the degree of Master of Arts

Department of Philosophy

2018

This thesis entitled:
A Defense of the Moral Error Theory
written by John Christmann
has been approved for the Department of Philosophy

_____

Graham Oddie

_____

Raul Saucedo

_____

Iskra Fileva

Date_____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline

Abstract

In the thesis, I provide a defense of the moral error theory. In the first section, I solve the formulation problem, which is the puzzle related to properly formulating the moral error theory without self-defeat. In the second section, I defend the claim that moral concepts commit their users to irreducibly normative reasons. Finally, in the third and last section, I defend the reduction argument against new criticisms and conclude that there are no irreducibly normative reasons.

Table of Contents

Chapter One

The Formulation Problem

1. The Formulation Problem

Moral error theory is a strange beast. It seems to be one of the only philosophical

theories where formulating it runs the risk of incoherence. On an initial attempt, you may

be tempted to formulate it like this:

MER1: All moral propositions are false.[1]

The formulation I have named **MER1** may seem to be fine, but it results in incoherence

if we assume that propositions can only either be true or false, and if a proposition is

false then its negation is true. If all moral propositions are false, then it seems to follow

that a moral proposition and its negation are both false. But we just assumed that the

negation of a false moral proposition is a true moral proposition. Surely the negation of

the false proposition, "It is not morally wrong to torture babies just for fun" is true,

namely, "It is not the case that it is not wrong to torture babies just for fun".

Maybe you're thinking that the error theory should just stick to non-negated moral

propositions and allow for negated moral propositions to be true. This seems to lose the

spirit of the error theory, and it becomes a form of moral realism, at least on one

influential conception of realism. There are two requirements for being a realist: first,

one must believe that that domain of discourse has to involve the expression of

propositions and second, at least some of those propositions are true (Sayre-McCord

1988). On his view, then, the amended formulation of the error theory actually qualifies

---

[1] I will refer to propositions from now on because I think that the falsehood of a proposition entails the falsehood of many of the other truth-bearers discussed in the literature, such as judgments, thoughts, beliefs, and sentences.

as a form of moral realism, which seems to be a problem since error theories are anti-realist. We need a solution to this problem, but we also need a name for this problem. Charles Pigden calls it The Doppelganger Problem, but I will stick with The Formulation Problem seems like a more accurate description of the issue at hand (Pigden 2007).

The Formulation Problem needs a solution before we continue on to discuss the argument for moral error theory. If we cannot even formulate the moral error theory without collapsing into incoherence or realism, then there is either no reason to motivate it with arguments or it just becomes an eccentric version of moral realism that none of the contemporary defenders of error theories would subscribe to. In section two, I will survey some of the existing attempts at solving The Formulation Problem and provide some reasons to doubt their plausibility. In section three, I will formulate my own solution to The Formulation Problem and defend it against potential criticisms. Only then can we sensibly move on to defend the error theory.

2. Possible Solutions to the Formulation Problem

There are several proposed solutions to The Formulation Problem in the contemporary literature. In this section, I will survey these potential solutions and provide reasons to doubt their plausibility. Let's begin with Charles Pigden's proposed solution. Pigden wants to save the idea that all moral propositions are false as the core of what he calls metaethical nihilism. However, he recognizes that this formulation results in the problems I mentioned in the previous section. He wants to formulate metaethical nihilism so that it avoids The Formulation Problem while it retains the spirit of the original problematic formulation.

2.1. Pigden's Solution

Pigden's first pass at a solution is his proposal that we restrict the scope of metaethical nihilism to non-negative atomic moral propositions. An atomic proposition is a proposition in which an n-place predicate is ascribed to n specific items (Pigden 2007). What makes an atomic proposition moral is when a moral predicate is ascribed to an item (or items). The moral predicates that Pigden is concerned with are those that pick out thin moral properties like goodness, rightness, and wrongness.[2] So an example of an atomic moral proposition is, "What Tim did to Anna was wrong". In this proposition, the predicate is, "wrong" and the item it is applied to is "what Tim did to Anna". A negative atomic moral proposition is an atomic proposition to which negation has been applied—it is the negation of an atomic moral proposition. An example of a negative atomic moral proposition is, "What Tim did to Anna was not wrong". In this case, the predicate ascribed to the item is, "not wrong" and the item is, "What Tim did to Anna".

The problem with Pigden's first attempted solution should be obvious. When we say that it is not the case that what Tim did to Anna is wrong, we are either saying that what Tim did to Anna is right, or that what Tim did to Anna is permissible. In either case, we are ascribing a non-negative predicate to an item. So, it seems as though negative atomic moral propositions sometimes entail or are logically equivalent to non-negative atomic moral propositions. The Formulation Problem reappears because it seems to be the case that negative atomic moral propositions sometimes either imply, or are

---

[2] I will only be discussing thin moral properties and predicates because it is unclear whether or not thick moral predicates have sufficient non-normative content to survive the loss of their normative content upon discovery of systematic normative error such that propositions involving them are not all false. Getting into that debate would take us too far afield, so I will table this discussion for another time.

equivalent to, non-negative atomic moral propositions. Similarly, to say that somebody is not tall is to say that they are either short or average height. To say that something is not wrong is to say that it is either right or permissible. So Pigden's proposed solution doesn't work.

Pigden attempts to solve the problem by rejecting the rules governing the transition[3] from some negative atomic moral propositions to non-negative atomic moral propositions. In essence, he is denying that, "X is not wrong" implies, "X is right or X is permissible". Unfortunately, Pigden fails to provide sufficient rationale for his rejection of transition rules. Rather, he appeals to the goal of preserving the spirit of metaethical nihilism as the motivating factor in rejecting the transition rules. However, this move may undermine the main argument for the error theory, namely the argument from queerness. The argument from queerness relies on a conceptual or semantic component that appeals to ordinary moral discourse and claims to find something metaphysically problematic about its ontological commitments. It seems as though the transition rules that Pigden wants to reject may be part of ordinary moral discourse. Since my aim in this thesis is to defend a form of the argument from queerness, I cannot endorse Pigden's rationale for rejecting the transition rules. Rather, I must show that ordinary moral discourse either does not include those transition rules, or it is ambiguous enough that rejecting the rules does not ignore the metaphysically problematic aspects of moral discourse. So, Pigden's solution is incomplete for my purposes.

2.2. Sinnott-Armstrong's Solution

---

[3] I am using Pigden's terminology here, which seems to be another way of saying, "rules of inference".

Another strategy that may avoid The Formulation Problem is to spell out the error theory in terms of positive first-order moral claims being uniformly false. Walter Sinnott-Armstrong attempts to sidestep The Formulation Problem this way, and in the process of doing so he defines moral error theory as the view that all positive first-order moral claims are uniformly false (Sinnott-Armstrong 2006, 34-36). Unlike Pigden, Sinnott-Armstrong only considers moral propositions about what moral agents ought to do or ought not to do, moral reasons, moral goodness and badness, etc. Mere moral permissibility does not play a role in this account (Olson 2010).[4] Because Sinnott-Armstrong's account leaves out mere moral permissibility, it fails to adequately account for ordinary moral discourse, which seems to imply the concept of moral permissibility and its cognates such as, "being allowed to do X". Perhaps Sinnott-Armstrong could respond by arguing that moral permissibility is not a real property that is picked out by the predicate, and since the error theory is about the nonexistence of moral properties, it follows that the propositions that turn out to be uniformly false are those that are supposed to refer to moral properties. So, propositions about moral permissibility fall outside of the scope of the propositions that are targeted by the error theorist.

For the moral property response to work, it would have to be the case that moral permissibility is not a real property. One way we could work out this response is by arguing that predicates about moral permissibility are actually veiled disjunctions about moral rightness or goodness or the action falling outside of the scope of moral

---

[4] If Sinnott-Armstrong intends for his account to cover ascriptions of permissibility and impermissibility as property ascriptions, then his view is the same as Pigden's and it suffers from the same problem as I mentioned in section 2.1. For the rest of this section I will set this possibility aside for the sake of avoiding repeating myself.

relevance.[5] However, unless the disjunction includes at least one non-moral disjunct, then it is arguable that moral permissibility is a moral property. To avoid this, one could either adopt the view that there is at least one non-moral disjunct or the view that there are no disjunctive properties. If we adopt the view that the disjunction involves at least one non-moral disjunct, then we need to identify what that disjunct is. Fortunately, I already listed a candidate above. "Falling outside of the scope of moral relevance" seems not to be a moral predicate. So, perhaps the error theorist could argue that, since there is a non-moral disjunct, the whole disjunction is not moral, so moral permissibility is not a moral predicate. I grant that this is a genuine possibility and I have no direct argument against it, but I will not adopt it because it still seems to me that it fails to respect ordinary moral discourse. When we use the predicate, "moral permissibility", we seem to be ascribing a property to actions or omissions.

This leads us to the next move that the proponent of this formulation of the error theory could make. Perhaps there are no disjunctive properties, so a disjunctive predicate is not a predicate that ascribes a property. When it comes to this possibility, I will avoid directly engaging with it as the debate about disjunctive properties will take us too far afield. What I will say is that there are plausible accounts of property individuation that entail the existence of disjunctive properties (Oddie 2005, ch. 6). Furthermore, the disjunctive interpretation, while defensible, seems to me to do violence to ordinary moral discourse. So, I will now move on to other solutions to The Formulation Problem and set this dialectic aside.

---

[5] The list is not meant to be exhaustive. The disjunction could be very long, hence the need for a shorter predicate like, "is morally permissible".

2.3. Systematic Falsehood

Perhaps we could formulate the error theory this way:

MER2: Moral discourse is systematically false.

This formulation of the error theory does not have the air of seeming incoherence like **MER1** did, so it has that going for it. However, with that benefit comes the cost of ambiguity and vagueness. What exactly is it for a discourse to be systematically false? I can wrap my head around all moral propositions being false, but moral discourse is not as the kind of entity that can be ascribed or denied a truth value. Furthermore, what is it for a discourse to be systematically false? Is there an exact number of propositions that have to be false for an entire discourse to fall within the scope of an error theory? Maybe there are just some moral propositions that form the core of moral discourse and if they turned out to be false, then the entirety of moral discourse would be rendered systematically false. This seems to be the most plausible way of spelling out this formulation. However, there are still problems with identifying those core moral propositions and with explaining how those propositions turning out to be false would render the entire discourse systematically false.

While **MER2** might seem like a plausible way of avoiding The Formulation Problem, it fails. The most promising set of moral propositions that can be targeted such that, if they turn out to be false, then moral discourse satisfies **MER2**, is the set including such propositions as, "X is wrong". But if, "X is wrong" is false, then we have enough for The Formulation Problem to come back into play. If, "X is wrong" is false, then, "X is not wrong" is true, and if, "X is not wrong" implies, "X is right or permissible",

then we have The Formulation Problem. **MER2** turns out to lack the resources to deal with the problem.

2.4. Truth-Value Gaps

One promising way of avoiding The Formulation Problem is to reject the principle of bivalence, which is the principle that there are only two truth values, true and false. The Formulation Problem seems to rely on bivalence, since it has to do with the shift from negative to non-negative atomic moral propositions, which appears to depend on The Law of Excluded middle. Rejecting bivalence involves positing at least one additional truth value or a truth-value gap.[6] For the sake of simplicity, I will assume that this strategy only requires one additional truth value, and I will also remain neutral on what it is, so let's call it untrue.[7] For a proposition to be untrue, it must be neither true nor false. So, the new formulation of the error theory says that all moral propositions are untrue. One way of motivating this formulation is to diagnose the erroneous aspect of moral discourse as a case of reference failure. If referring expressions and propositions fail to refer, then perhaps those truth-bearers are neither true nor false. On this view, for a proposition to be false, it must still refer to something, but it falsely ascribes a predicate to that thing. In other words, it must say something false about something that exists.

---

[6] For example, if propositions are functions from world-times to truth values, then some propositions are not complete functions, so they yield no truth value when applied to a world-time. A useful analogy is the division function. Applying the division function to the pair, [1,0] yields nothing. In other words, dividing one by zero gets you nothing. In the same way, it's possible for propositions to yield no truth value rather than a third one. I am indebted to Graham Oddie for pointing this possibility out to me.

[7] The truth value gap mirrors what I call the permissibility gap in deontic logic. The truth value gap shows up in cases where a proposition is neither true nor false, so it falls within the scope of, 'untrue'. Being untrue could be either being false or being neither true nor false. In deontic logic, it is usually assumed that situations are either permissible or impermissible (to be obligatory implies permissibility, $(O(x) \rightarrow P(x))$, because you cannot be obligated to do something if that action is impermissible). But, if my preferred solution to The Formulation Problem is correct, then there must be permissibility gaps where actions are neither permissible nor impermissible. According to my view, the negation of permissibility could either imply impermissibility or falling outside the scope of the permissible/impermissible dichotomy.

The way in which this proposal could sidestep The Formulation Problem is by providing a way to reject the transition rules that allow us to move from some negative atomic moral propositions to non-negative atomic moral propositions. Recall that the problem arises because propositions like, "X is not wrong" seem to imply propositions like, "X is either right or permissible". The reason is that negation seems to work. A negated atomic moral proposition seems to entail a non-negative atomic moral proposition based on rules governing the shift from negated moral propositions to moral propositions ascribing moral predicates to items. This appears to be a promising strategy for an error theorist to adopt when confronted with The Formulation Problem. However, I will not endorse this strategy because I am skeptical about truth-value gaps, and I am suspicious that the view confuses property ascription with reference failure due to the non-existence of objects being referred to. Before I move on to discuss my own response to The Formulation Problem, I will briefly motivate my suspicions and thus motivate my hesitant attitude towards adopting the truth-value gap strategy.

In the famous example, "The present King of France is bald", the part of the sentence that aims at referring to the present King of France fails to refer, which is why the sentence as a whole could arguably lack a truth-value, although that is also debatable. When it comes to moral sentences or propositions, the part that refers to an object picks out an agent, event, or action, and in that sense, it does not fail to refer. Moral sentences or propositions contain moral predicates that aim at ascribing properties to those agents, events, or actions, however, those properties do not exist if the error theory is true, so moral sentences or propositions ascribe non-existent properties to existing or possible objects, and those sentences or propositions are false.

For example, if I say, "My car is red all over and green all over" I have said something false because I have ascribed a non-existent property to an existing object. So, sentences or propositions that ascribe non-existent properties to existing or possible objects are false, not some third truth value or lacking a truth value (Streumer 2017, 123-124). The proponent of this strategy has confused the failure of referring expressions picking out objects with predicates ascribing non-existent properties to existing objects.[8] Before I move on to my solution, I want to note that this seems to be a viable strategy for error theorists who are more open to allowing truth-value gaps into their ontologies than I am. If I were to be convinced that my solution fails, I would find myself siding with the truth-value gap crowd, so I want to make sure that I am not multiplying my enemies, but rather flagging a worry about a possible solution to The Formulation Problem.[9]

2.5. My Preferred Solution

My preferred solution to The Formulation Problem is like Pigden's. I think that moral propositions contain predicates that ascribe properties.[10] Those predicates include, "permissible" and "impermissible". Those predicates do not logically entail the other when negated, so, "not permissible" does not entail, "impermissible" and, "not impermissible" does not entail, "permissible". The appearance of such logical entailment

---

[8] One possible response that was brought to my attention by Graham Oddie goes like this: The so-called "reference failure" could turn up in the predicate position. Consider Q: Trump instantiates the present King of France's favorite property. Since there is no King of France there is no such property as the present King of France's favorite property. Arguably Q is neither true nor false. My response is that Q seems to be false, because we are saying something false about Trump. We are falsely saying that Trump instantiates some property that does not exist, so Trump couldn't actually instantiate such a property, so Q is false.

[9] This section is heavily influenced by several conversations I had with Graham Oddie, an avid supporter of truth-value gaps, and to whom I owe a great deal.

[10] I will not take a stand on what the bearers of those properties are supposed to be. Possible candidates include states of affairs, events, persons, and actions.

is created by conventions or substantive assumptions employed in ordinary moral discourse that connect negated moral predicates with other moral predicates. The conventions governing the transition between negated moral predicates and other moral predicates can be explained in terms of conversational implicature.[11] Such implicature can be cancelled in contexts like the seminar or conference room where the error theory is being articulated and discussed. The act of articulating the error theory generates a conversational context in which it is no longer permissible to shift from negated moral predicates to other moral predicates. So, it appears we can preserve bivalence by denying that the error theorist must allow for transitions like, "not permissible" means or entails "impermissible", and, "not impermissible" means or entails "permissible".

One reason to think that my preferred solution is correct is to compare permissibility with tensed predicates. For example, the sentence, "Prime numbers are morally wrong" is false, which means that, "Prime numbers are not morally wrong" is true, but that does not entail that prime numbers are permissible. The sentence, "Prime numbers are permissible" doesn't seem to make sense because prime numbers aren't

---

[11] I will not go in depth into what kind of conversational implicature is at play in the part of moral discourse involving negated predicates. That is a topic for another occasion, but for now it is sufficient to point out that the three mainstream theories of conversational implicature are all capable of providing plausible explanations of the phenomenon of shifting from negated moral predicates to other moral predicates. For an articulation of the relevance theory of conversational implicature, see (Sperber and Wilson 1981). The relevance theory can identify contexts of conversational relevance at play in ordinary discourse that are not present in situations where somebody articulates an error theory. Also see (Horn 2004) for a Neo-Gricean theory whose R principle says that we should say no more than we must, which allows for the context of conversation to determine how much and what must be said to convey speaker-meaning. In the case of ordinary discourse, speakers assume that the other person they are talking to are working within moral discourse such that they do not need to explain how they are ascribing moral properties to property-bearers. When discussing the error theory, such ascription is made explicit. For a classic view of conversational implicature, see (Grice 1957). On Grice's view, there is a maxim of relation which tells us to be relevant, which means that in ordinary moral discourse it would be inappropriate not to transition from, "is not permissible" to, "is impermissible, since ordinary moral discourse users probably aren't error theorists, and most error theorists outside of the seminar or conference room are not acting as such (cf. Grice 1957, 26-30). So, it seems to me that on every mainstream view, the error theorist can explain the conversational implicature governing the transition from negated moral predicates like, "is not permissible" to predicates like, "is impermissible". For another view about conversational implicature and the error theory, see (Olson 2014).

the sorts of things that can bear the property of permissibility, which seems to imply that the sentence involves a category mistake. This is not an *ad hoc* consideration, either, and that's where tensed predicates come in. We can plausibly say, "Prime numbers are late"[12] which is false, implying that prime numbers are not late, but prime numbers are not thereby early. Prime numbers just don't bear those sorts of properties. My solution to The Formulation Problem explains this phenomenon. In the conversational contexts where discussion of the metaphysics of numbers arises, sentences like, "Prime numbers are not morally wrong" will be about prime numbers failing to bear the property of moral wrongness without entailing or implying that numbers bear the property of permissibility. Further, in contexts about the relationship between time and the nature of numbers, it may be sensible to say that prime numbers are never late, meaning that prime numbers never bear the property of being late, and all of this can happen without entailing or implying that prime numbers are always early or on time. Similarly, in contexts where the discussion of actions and other moral property bearers arise, we can express propositions about actions not bearing the property of permissibility with sentences like, "Doing X is not permissible", without implying that doing X is impermissible.[13]

3. Kaspar's Objection and My Reply

The main objection to my articulation of the error theory comes from David Kaspar. He argues that, since error theorists commit themselves to an intuitionist semantics, they

---

[12] I am assuming that the conversational context will imply a time at which the lateness of something is indexed to, and that relations are genuine properties.

[13] This way of motivating my preferred solution to The Formulation Problem was suggested to me by Bart Streumer, to whom I owe a great deal from our discussions about the issues discussed in this section.

cannot introduce notions like conversational implicature because those are additions to

the semantics that conflict with what intuitionists have formulated (Kaspar 2017). Since

intuitionist semantics requires that there is a logical entailment between, "not

permissible" and, "impermissible", any talk of conversational implicature is beside the

point, since it would violate the entailment between those two predicates.[14] In other

words, intuitionist semantics says that, "not permissible" entails, "impermissible" and

any talk of conversational implicature determining the appropriateness or

inappropriateness of such inferences is to depart from intuitionist semantics and to

change the subject. So, Kaspar's argument is that error theorists are committed to

intuitionist semantics, and since intuitionist semantics is governed by The Law of

Excluded Middle, propositions with the negated moral predicate, "not permissible" entail

propositions with the moral predicate, "impermissible". Error theorists cannot coherently

articulate their view because a false moral proposition entails a true moral proposition,

so it cannot be the case that all moral propositions are false. The negation of a moral

proposition is another moral proposition, so error theorists cannot be simultaneously

committed to the claim that all moral propositions are false, and intuitionist moral

semantics.[15] That seems to be the entirety of Kaspar's argument, and I see two

problems with it. First, he does not provide any evidence that intuitionists have explicitly

articulated their semantics as including logical entailments between negated moral

predicates and other moral predicates. Second, error theorists are not necessarily

committed to intuitionist semantics. I will unpack these responses in order.

---

[14] Kaspar is talking about metaethical intuitionism, which is a view that is closely related to metaethical non-naturalism, and not intuitionist logic.

[15] I want to reiterate that I am discussing metaethical intuitionist semantics about moral predicates and not intuitionist logic, in case of any confusion that might arise because of the terminology.

First, Kaspar fails to provide evidence that the consensus view on moral semantics among intuitionists includes logical entailments between negated moral predicates and other moral predicates. While this objection does not show that Kaspar is wrong, it does show that he must do more work to convince us that he is right. In my own experience of reading intuitionists, non-naturalist and naturalist variants included, I have not encountered any explicit acknowledgement of such logical rules, and while they may be implicit in the writings of intuitionists, I have failed to find such an implicit consensus. But I may be wrong, which is something Kaspar must show, and the best way to do that is by providing evidence of an explicit or implicit consensus about such logical rules.

Second, Kaspar only provides one instance of an error theorist committing himself to intuitionist semantics, namely Jonas Olson (Olson 2014, 14). Error theorists are free to articulate their theories as they see fit, but it is not clear that from Olson's view we can generalize to all error theorists. For example, Don Loeb's error theory, which he calls moral incoherentism, is explicitly not committed to intuitionist semantics. Loeb believes that we may empirically discover that moral discourse involves conflicting semantics, the co-presence of which within one discourse generates incoherence, and thus an error theory about such discourse (Loeb 2008).

I will articulate my own view about the semantic commitments of ordinary moral discourse in the next section, but I will briefly explain how they are not strictly in line with Kaspar's intuitionist semantics. I believe that we must use a combination of empirical inquiry and considered judgments about particular cases to come to a view about moral semantics. I will argue that, ultimately, our moral semantics commit us to irreducibly

normative properties. However, I think that we must hold to a cautious optimism about those commitments because we have not engaged in sufficient empirical inquiry investigating whether ordinary moral semantics commits its users to the existence of irreducibly normative properties. The caution is due to the lack of empirical evidence going one way or another, but the optimism is based on what we can gather from considering particular cases and looking at how we, ordinary users of moral discourse, talk. I think that we must look at the function of moral discourse within contemporary societies to ascertain what we commit ourselves to when we use it as ordinary speakers. So, while I do think that ordinary moral discourse ontologically commits its sincere users to irreducibly normative entities, much like many non-naturalist intuitionists, I do not think that there is enough empirical evidence to adjudicate between the conversational implicature claim I made in the previous section and Kaspar's claim about intuitionist semantics when it comes to which one reflects ordinary usage more accurately.[16] In the meantime, since there seem to be good reasons for cautious optimism about the claim that ordinary moral discourse commits us to irreducibly normative entities, and since I will present and defend an argument against the existence of such properties, it seems worthwhile to continue this thesis given that error theory as I articulated it is a live option.

---

[16] However, as I will argue for in the next part of this thesis, I take philosophical analysis to be first-personal and egoistic, so a successful analysis need not rely on empirical evidence of the kind that is unavailable to philosophers sitting at their desks. In other words, philosophers should probe their own use of terms through the tools available to them like thought experiments, and not worry so much about how others speak. In that regard, I do not consider empirical evidence to be all that relevant, but it is a nice way to bolster my case if it turns out that lots of other people use moral concepts like I do, and it makes my theory more interesting because it has a much wider scope.

Chapter Two

The Function of Moral Concepts

1. Preliminaries

Contemporary moral discourse commits its users to irreducibly normative entities.[17] Let

me unpack that a bit before proceeding. My claim is about what speakers commit

themselves to by using moral discourse. By moral discourse I mean deploying moral

concepts to make sincere moral judgments either in one's head, by speaking, or by

writing them down.[18] What moral discourse commits its user to is determined by the

functions of the moral concepts being deployed in making sincere moral judgments. I

will argue for a cognitivist account of the function of moral concepts according to which

their function is to match certain kinds of entities to other kinds of entities such that

sincere moral judgments that deploy those concepts are truth-apt. Their truth-conditions

are set by the matching function. For example, if I sincerely judge that you have acted

wrongly, I deploy concepts in my judgment, the function of which is to match the

property of moral wrongness to your actions.

The sorts of entities that our moral concepts match to descriptive entities are

probably quite varied. I do not take a stand on exactly what the entities are, nor do I

argue for any specific view on what the normative entities are, besides the claim that

whatever they are, they imply irreducibly normative reasons. Such reasons are

categorical insofar as they apply regardless of the desires, attitudes, or roles of the

---

[17] I do not take a stand on what those entities are just yet. They may be facts, properties, tropes, states of affairs, character traits such as virtues, judgments, etc. The list goes on. In the next section I will commit myself to using properties as an example of irreducibly normative entities, but the argument I defend there does not need the entities to be properties.

[18] There may be other ways of expressing sincere moral judgments that I've missed, but it seems to me that the categories I listed exhaust the things that we do when we engage in moral discourse.

person being criticized or advised. My strategy will be to identify the function of moral concepts in contemporary moral discourse by pumping intuitions about particular cases, and by examining the available empirical evidence about how ordinary speakers use moral discourse.

In what follows, I will show that moral concepts commit their users to irreducibly normative entities. In the next section, I will further articulate and motivate the metaphilosophy underpinning my methodology in this section. Then, in section three, I will defend a view that I call cognitivist descriptivism. In section four, I will explain what it is about moral discourse that commits us to irreducibly normative entities. In section five, I will examine some intuition pumps in the hopes that our considered judgments will vindicate my claim about irreducible normativity. Finally, in section six, I will wrap everything up.

2. Fumertonian Metaphilosophy

I hold a conciliatory view when it comes to the issue of a priori versus a posteriori semantic methodology. Philosophers should take note of how ordinary moral concept users speak seriously when analyzing moral semantics, but that is not the only relevant evidence available. Since philosophers live among ordinary moral concept users in their daily lives, especially when they are outside of academic settings, they can rely on their own linguistic dispositions as evidence of how moral concepts are used.[19] In addition to

---

[19] At least they can note the usage among their in-groups and maybe within their broader society. It could be that moral concepts are used differently by different linguistic communities (cf. Flanagan 2016). If it so happens that some linguistic communities use moral concepts in ways that do not imply an error theory on the grounds argued for in this thesis, then within those communities, it would be false that they are saying false things when they use moral concepts in making moral judgments. Whether or not relativism about moral concepts is true is outside of the scope of this thesis, so I will just say that if it's true, then the

this weak view of armchair philosophy, I will also briefly motivate a stronger view which says that philosophy is ultimately egoistic or first-personal. This implies that in doing philosophical analysis of the meanings of concepts or terms, a philosopher should rely on her intuitions as sources of insight into her linguistic dispositions. How she would use concepts across modal space is what she should me aiming to decipher. In what follows, I will further articulate and motivate this view.

It should be noted that evidence of how others use concepts is also of interest to a philosopher, because there are interesting issues that arise from considerations about mismatches between her own use and the way others use the concepts being analyzed. It would also be good evidence in favor of a proposed analysis if the philosopher could show that others have similar linguistic dispositions as she does. So, empirical evidence still matters, but it should not be the most important consideration when doing philosophical analysis of concepts. Empirical evidence and evidence from the egoistic, first-personal perspective can be mutually supporting, but in conflicts, the latter takes precedence over the former. Now I will briefly motivate the primacy of the egoistic, first-personal methodological perspective.[20]

I want to motivate a traditional, a priori perspective on doing moral semantics that borrows heavily from Richard Fumerton's paper, "The Paradox of Analysis". The view is that moral semantics can be done from the armchair because philosophical analysis of moral concepts is ultimately an exercise in discerning our linguistic dispositions. In that respect, philosophical analysis is egoistic or first-personal. When we give an account of

---

way the society in which I am embedded uses moral concepts implies an error theory, but maybe not in other societies. This seems to be an interesting enough conclusion.

[20] I should note that this blend of armchair and experimental methodologies has been adopted in other areas, such as the philosophy of race (Glasgow 2008) (Pierce 2015).

what we mean by particular moral phrases, we test it against possible counterexamples

that represent scenarios in which the application conditions for those phrases do not

match the proposed account. If those scenarios are genuine counterexamples, then the

proposed account fails. This methodology is widely used in philosophy, with the

paradigm being Gettier counterexamples to justified, true belief analyses of

KNOWLEDGE[21] (Gettier 1963). The methodological structure is simple: find a particular

case that matches the proposed analysis, and then show how the application conditions

in that scenario do not match the conditions of the analysis. In the case of Gettier's

examples, the cases match the conditions for a justified, true belief account of the

application conditions for KNOWLEDGE, but it is clear that those do not constitute

cases in which it is appropriate to apply KNOWLEDGE when making epistemic

judgments about the subject being discussed. The methodological structure is

instantiated by various other debates in philosophy, such as Hilary Putnam's use of the

twin earth thought experiment to refute internalism about meaning, Carl Cohen's new

evil demon argument against epistemic externalism, and Harry Frankfurt's

counterexamples to the principle of alternative possibilities.[22]

The considerations that I raised in the paragraph above may not be enough to

motivate the view for some people. Before finishing this section, I want to articulate a

thought experiment to elicit your intuitions about this topic.[23] Imagine that you become

---

[21] I use all caps when mentioning a concept.
[22] For twin earth, see (Putnam 1973), for the new evil demon see (Cohen 1984), and for the Frankfurt example see (Frankfurt 1969). For a compilation of famous thought experiments, see (Tittle 2004).
[23] This thought experiment is briefly gestured at by Richard Fumerton in his illuminating paper, "The Paradox of Analysis" (Fumerton 1983). I will try my best to flesh it out while remaining true to its Fumertonian character.

convinced that solipsism is true[24] because you spent an inordinate amount of time reading Descartes. But prior to spending so much time reading Descartes and updating your credences, you spent a lot of time analyzing the meaning of, "normative reason". You have a well thought out analysis of the notion of a normative reason, and you even wrote some papers and a book about it. Now that you are a convinced solipsist, should you also lower your credences when it comes to your well thought out analysis of the notion of a normative reason? It seems to me that you can remain confident in your analysis, as long as the reasons why you became a solipsist do not imply the falsehood of anything central to your analysis. You can remain confident that your analysis of the notion of a normative reason survives your conversion to solipsism because your analysis was ultimately about your linguistic dispositions concerning your use of, NORMATIVE REASON. The same seems to go for any analysis you might have engaged in, as long as the reasons you used to arrive at that analysis are not defeated by the reasons you deployed when becoming a solipsist.

While the thought experiment I articulated above may seem to commit me to a solipsistic account of philosophical analysis, I maintain that that is not the case. I do not discount the value of considering others' linguistic dispositions when doing analysis, especially when we are operating under the assumptions that: a) philosophical skepticism is false (that is, we have access to facts about the external world), and b) our linguistic dispositions aren't *too* idiosyncratic. But the truth of an analysis does not ultimately hinge on those assumptions. It is possible for those assumptions to be false, while an analysis is true. Such an attitude towards philosophical analysis is not entirely idiosyncratic, either. For example, when discussing the meaning of the word, "good" and

---

[24] By, "solipsism" I mean the view that we cannot know that there are other minds.

how he will analyze it, G.E. Moore distinguishes how he uses it from how most people use it, and says,

> "What, then, is good? How is good to be defined? Now, it may be thought that this is a verbal question. A definition does indeed often mean the expressing of one word's meaning in other words. But this is not the sort of definition I am asking for. Such a definition can never be of ultimate importance in any study except lexicography. If I wanted that kind of definition I should have to consider in the first place how people generally use the word, 'good'; but my business is not with its proper usage, as established by custom. I should, indeed, be foolish, if I tried to use it for something which it did not usually denote; if, for instance, I were to announce that, whenever I used the word 'good', I must be understood to be thinking of that object which is usually denoted by the word 'table'. I shall, therefore, use the word in the same sense in which I think it is ordinarily used; but at the same time I am not anxious to discuss whether I am right in thinking that it is so used" (Moore 1903, 6).

Moore's sentiment is shared by other philosophers (e.g. Marks 2013, 2-3). I take philosophy to be a very personal endeavor, where, at the end of the day it's about how convinced you, not others, are about a particular view or argument.[25] So, when it comes to the analysis of moral concepts, I will lean heavily on the Fumertonian metaphilosophy that I have articulated above.

3. Against Expressivism

---

[25] This may be a source of some intuitions that favor internalist analyses in epistemology as well. For example, it could be that an egoistic perspective about metaphilosophy informs how one theorizes about the epistemic, or vice versa.

In this section, I defend cognitivist descriptivism and attack expressivism and other opponents of descriptivism. I will establish a prima facie case for descriptivism and then I will show how denying descriptivism has theoretical costs that we should avoid. Cognitivism is the view that when we make moral judgments we express beliefs. So, cognitivism concerns moral psychology, or what we do when we make moral judgments. According to cognitivism, to make a moral judgment is to express a belief with certain contents. Descriptivism is the view that moral beliefs are representational, and moral beliefs aim to represent the world. In other words, the contents of moral beliefs picture or mirror moral reality correctly or incorrectly. If they do so correctly, they are true, and if they do so incorrectly, they are false.[26] Cognitivist descriptivism is the view that when we make moral judgments we express beliefs that aim to represent moral reality such that those beliefs can be true or false. In what follows, I will defend cognitivist descriptivism and expose the problems that come with denying it.

The fact that expressivists and other non-cognitivists have tried for decades to accommodate belief talk, reference talk, truth talk, fact talk, and property ascription talk seems to constitute evidence that moral discourse appears to behave according to the way semantic descriptivists predict. Non-descriptivist and non-cognitivist theories are constructed such that they accommodate those phenomena, but they rarely satisfactorily explain them in the unified and powerful way that cognitivist descriptivism does. Cognitivist descriptivism does not have the same problem, because the only areas where moral discourse seems to be in tension with descriptivist semantics involve imperatives like commands, and the motivational heat of moral judgments. But the

---

[26] I will not get into whether truth and falsity are categorical or on a spectrum based on truth-likeness. If they're a spectrum, then things can be more or less truth-like, meaning that the closer a proposition represents reality, the more truth-like it is (cf. Oddie 2014).

semantics of moral commands is a very controversial issue, and there are plausible truth-conditional semantics of commands (Parsons 2011). Moral imperatives are not as clear cut as judgments that take the form, "I believe that it is true that the act of deliberately harming an infant only for your own pleasure instantiates the property of wrongness or impermissibility." When it comes to moral judgments like the one I just mentioned, it is quite clear that it is well formed and our ability to express and understand it is better explained by cognitivist descriptivism. So, the areas of moral discourse that seem to be in tension with cognitivist descriptivism are minimal and not clearly as problematic as the areas in tension with non-descriptivism and non-cognitivism.

The motivational heat of moral judgments is not a semantic issue, so it falls outside of the explanatory scope of descriptivism. Issues of motivation are within the purview of moral psychology, not semantics. So, descriptivism seems to have a prima facie upper hand, and expressivists have been trying to accommodate things that naturally hang together on a descriptivist theory of moral semantics. But it might impugn cognitivism, which means it would still be a problem for cognitivist expressivism. My response to worries about cognitivism being incapable of accounting for motivational heat is that many instances of moral judgments do not seem to involve any sort of motivational heat that expressivists talk about. For example, many people are fully aware that they can donate a nominal amount of money to charities that can significantly improve, or even save, the lives of those much worse off than anybody in the developed world, and those people will even say that we ought to donate to such

charities, but those people just don't do it.[27] Those cases seem to me to be instances of

sincere moral judgments that aren't motivationally hot.

Another set of examples involve judgments about the past. We can make sincere

moral judgments about the wrongness of past events or the badness of past historical

actors without the phenomenology associated with motivational heat arising. When I

make a judgment about how wrong it was for ancient Athenians to recognize and

endorse a slave class within their society, I do not seem to have the sorts of

experiences associated with being motivated to act on a judgment I make. Unlike cases

where I judge that I ought to help my immediate friends or family, or that something

done to them was morally wrong, cases involving very geographically, or historically

distant events or actors do not seem to reliably produce the phenomena that indicate

motivation. Considerations like these are corroborated by recent work on moral

psychology. While the so-called, "affective-revolution" has put moral cognitivists[28] on the

defensive, the evidence they produce does not extend to cases like the ones I

mentioned; at best, the empirical evidence produced by expressivists is restricted to

judgments about the subject's immediate friends, family, and others within their

proximity, both in time and space (cf. McAuliffe MS).[29] In other words, the closer the

situation is to the self of the person who makes a moral judgment, the more likely that

---

[27] One reply may be that those people fail to make genuine or sincere moral judgments, because some variant of internalism is true. Addressing arguments for or against motivational internalism about moral judgments would take us too far afield but suffice it to say that adding internalist machinery seems to be an attempt at addressing a defeater by complicating the overall theory. My strategy of denying that all genuine or sincere moral judgments are motivationally hot seems to be an overall simpler move in light of the counterexamples I raise.

[28] At least those cognitivists who are not also sentimentalists.

[29] I must also note that a lot of the evidence produced by recent work in moral psychology is compatible with a hybrid view that takes moral judgments to be mixed entities that have a non-cognitive, motivational part, and a cognitive, representational part. Such states may be called, "besires" which is a mix of belief and desire. For a useful review of the ways a besire theory can be developed, see (Edwards 2011). .

the judgment will be accompanied by an affective state like an emotion. The further from the self the situation being evaluated is, the less likely that the judgment will be accompanied by an affective state. For example, most people would not experience strong emotions when making the judgment that killing Julius Caesar was unjustified or justified, but they will probably experience such emotions when they make a judgment that the slaying of a family or in-group member was justified or unjustified. Since expressivists and other non-cognitivists take moral judgments to be constituted or identical to affective states[30] like emotions or pro-attitudes, these cases present counterexamples to their analyses of moral judgments.

So, it seems to me that there are prima facie reasons to reject expressivism in favor or some form of cognitivism. While there may eventually be overridingly strong evidence in favor of expressivism that can overturn the prima facie reasons, when we look at moral discourse, it just seems like cognitivist descriptivism captures important elements of it. The existing evidence used to put such views on the defensive does not actually support expressivism, while it may, at best, support some sort of mixed theory with core cognitivist-descriptivist elements.[31] I conclude this section with cautious optimism about the prospects of cognitivist descriptivism.

4. The Functions of Moral Concepts

In this section, I will explain how ordinary moral discourse commits its sincere users to irreducibly normative entities. I think that sincere speakers using moral discourse

---

[30] For an overview of contemporary non-cognitivist theories, see (Miller 2003) and (Schroeder 2010).
[31] The existing evidence does not strongly support the view that moral judgments are entirely constituted by mental states that have built in motivationally hot parts, like desires or commitments to plans of action, which is what the expressivist needs to get her theory off the ground with empirical evidence.

commit themselves to irreducibly normative entities through the deployment of moral concepts. Typical moral concepts include: MORALLY PERMISSIBLE[32], MORALLY WRONG, and MORALLY BLAMEWORTHY. Such concepts match irreducibly normative entities to natural entities. For example, the concept, MORALLY WRONG can match the normative property of moral wrongness to the actions of a particular person.[33]

If we look to the way that we use moral discourse in contemporary society, then we can discern the central functions of a moral concept,[34] including matching irreducibly normative entities to non-normative entities.[35] [36] In other words, if we look to the function of moral discourse in our current society, we can determine how moral concepts function, and how they function is a window into what sincere speakers commit themselves to through their use. My contention is that moral concepts have a clear and essential function, which is to guide action.[37] [38] Moral concepts guide action because they are either about or they imply reasons for action. Some moral concepts

---

[32] I use all caps when mentioning concepts.

[33] I am assuming that actions done by moral agents are describable in terms of natural properties or something like that.

[34] These are functions that are essential to moral concepts. This means that if moral concepts ceased to perform those functions, they would no longer be moral concepts. I believe that this is the best way to articulate what Richard Joyce means by non-negotiable commitments of moral discourse (cf. Joyce 2001).

[35] When I say it matches entities to entities, I mean that it aims to represent something like the instantiation of some entities by other entities. For example, moral concepts that appear in predicate expressions represent the instantiation of certain moral properties by certain moral property-bearers, like, "is morally wrong" when connected with, "What Tim did" represents the instantiation of moral wrongness by whatever it is that Tim did. I do not want to get too deep into the mechanics behind this, because it would take us too far off topic, but I will address the issue of cognitivism vs. non-cognitivism further on in this section of the thesis. I am just giving an account of how, on a cognitivist view of moral discourse, the use of moral concepts for things like predicate expressions can commit sincere speakers to irreducibly normative entities. When it comes to the matching relation, I intend for it to be able to be a representational stand in for relations like instantiation, exemplification, and others like it.

[36] I want to note that it could be the case that normative properties also have normative properties, so the fact that I did the right thing could have the property of being good. I leave this complication aside and stick with talk of matching normative properties onto descriptive properties.

[37] Essential functions are functions that either disclose or constitute the identity conditions of a concept. I remain neutral about whether concepts' identity conditions are determined by their functional roles or something else.

[38] This isn't their only essential function, but it's all I need for my argument in section three to work.

like, MORAL REASON are obvious cases of concepts that guide action through their application. Other concepts like, MORALLY WRONG don't wear their action guidance credentials on their sleeves, but they still imply reasons to act. If I tell somebody that their action is morally wrong, I am implying that they ought not to engage in that action, or, in other words, they have a reason to refrain from acting that way.

But why think that moral concepts are intimately connected with the concept of a reason for action? Besides the oft repeated slogan that ethics and morality ought to be action guiding, what reasons are there to think that reasons and morals are related in a non-negotiable way? Consider how we use moral discourse when we criticize the behavior of other moral actors. It seems clear that from the moral point of view, an actor ought to correct her behavior in light of moral considerations.[39] When we say that someone ought to correct her behavior in light of moral considerations, that seems to be another way of expressing the proposition that she has reasons to correct her behavior. When I tell somebody that what she did was wrong and that she ought not to do that again, I am suggesting that she has a reason not to do thing again. If, in response to my moral evaluation, she responded with, "Sure, what I did was wrong, and I ought not to do it again, but there is no reason for me not to do it again", I imagine that we would consider her confused or playing some sort of game. Such utterances strike me as close to Moorean paradoxes; in both cases, there is a tension if not outright contradiction between the two parts of the utterance. In the case of Moore's paradox, the tension arises between the actor uttering that P is true but that she does not believe that P. In my case, the tension arises between the actor uttering that what she did was

---

[39] I do not take a stand on if she ought to be aware of the salience of those moral considerations. My argument works either way.

wrong, and she ought not to do it again, but that there is no reason for her to refrain from doing it again.

What does the Moorean-like phenomenon I mentioned above tell us about the function of moral concepts? I think that it shows that there is an intimate connection between the concept of a reason for action and moral concepts like, WRONG and OUGHT/OUGHT NOT. The sort of connection I envisage is similar to the one revealed by Moore's paradox between truth and belief. The concept of truth is, in some sense, built into the concept of a belief; beliefs without the goal of truth seem empty or just not beliefs at all. While I do not want to suggest that there is an action goal analogous to the truth goal for beliefs, I see a tight link between the concept of a reason for action and moral concepts. Without that link, much like with beliefs and truth, moral concepts would seem empty or pointless. They may even lose their status as moral concepts because their identity conditions heavily rely on reasons for action.[40] But, you may insist that I am stacking the deck in my favor by only dealing with deontic moral concepts like, WRONG and OUGHT. I think the same Moorean phenomenon arises when we consider evaluative moral concepts like, GOOD and BAD. For example, it seems just as odd to say, "I see that what I am doing is bad, but there is no reason for me to stop doing it". It also seems odd to say, "I see that what I am doing is good, but there is no reason for me to continue doing it." Evaluative moral concepts seem to exhibit the same properties in this respect.

---

[40] I will not venture down the rabbit hole of what exactly the identity conditions of concepts are. I am sympathetic to an inferential role semantics of concepts, so their identity conditions are determined by their functional role in inference, but my argument is neutral when it comes to what determines conceptual identity conditions.

Perhaps, you insist, I am still stacking the deck because I keep mentioning actions in the sentences I mentioned above. I think the same phenomenon also arises when we consider moral evaluations of behavior and character that do not directly mention actions or omissions. For example, imagine that upon being informed that an actor has cultivated the vice of envy. She is told that she regularly exhibits behaviors that somebody who had an envious character would engage in. Upon hearing this, she says, "Well, sure it's bad to be envious, but there is no reason to cease being envious." Again, the phenomenon occurs because she acknowledges that being envious is bad, but then she follows up by saying that there is no reason to cease being envious. It seems like the badness of envy generates or implies reasons to avoid being envious in the future. Let's look at other evaluations. Perhaps an actor is told that consuming a particular mushroom harm her, and that the harm she would incur is bad. She then says, "I realize that consuming that mushroom would harm me and that sort of harm is bad, but there is no reason for me not to consume that mushroom." In this case, the phenomenon still arises. Now maybe something fishy is going on because we are discussing harms and vices had by the actor being evaluated. Maybe the phenomenon arises because she is the subject of evaluation, and if we evaluated somebody besides her, when she uttered those sentences she wouldn't sound as strange. Imagine that the actor and I see somebody who has consistently acted enviously. I tell my companion that this person is exhibiting the vice of envy and envy is bad. My companion responds with, "Yes, she is certainly an envious person, but there is no reason for her to cease being envious." The strangeness persists. Now, let's imagine the mushroom scenario, but the actor and I see somebody else about to consume the harmful mushroom. I tell

my companion that the mushroom is harmful, and that harm is bad. My companion

responds, "Sure, the mushroom will harm her and that is bad, but there is no reason for

her not to consume the mushroom." There seems to be no difference between

evaluating the person uttering the strange sentence and evaluating somebody else.[41]

So, there seems to be a persistent phenomenon when it comes to the concept of

reasons and moral concepts that in some sense mirrors the Moorean paradox

concerning truth and belief. I take this to be strong evidence of an intimate connection

between the concept of a reason for action and moral concepts, both evaluative and

deontic. But what are reasons and why think that they are irreducibly normative? After

all, this whole excursion into moral concepts is about trying to show how moral

discourse commits its users to irreducibly normative entities. Let's take a look at the

metaphysics of reasons.

5. The Metaphysics of Reasons

What sorts of entities are reasons? Why are they irreducibly normative? I take reasons

to be a kind of relation. When something is a reason for something else, there is a

relation between those two things that we can call the reason relation. Feeling thirsty is

a reason to drink water, which means that there is a reason relation between feeling

thirsty and drinking water. We can represent the reason relation with a formalism to

---

[41] Perhaps there are sentences like, "The fertilizer is good for that plant, but there is no reason to fertilize that plant" which may not strike us as strange. First, it still seems strange to me. Even if the plant was an invasive species or a weed, there still seems to be a pro tanto or prima facie reason to fertilize it because the fertilizer is good for it, but that reason can be overridden by other considerations, meaning there is no all things considered reason to fertilize. Second, even if it isn't strange, it isn't clearly a moral sentence. Rather, it seems to be a non-moral, evaluative sentence about plant welfare, which does not appear to be within the domain of morality. My discussion is about moral concepts, not evaluative concepts as such.

abstract away from the particularities of the things that stand to each other in that relation.

Reason Relation: R(F, A, Φ, C).[42]

We can read this as F is a reason for A's doing Φ in circumstance C. But what exactly is irreducibly normative about this relation? Things that can take the place of F, A, Φ, and C do not seem to be irreducibly normative. Being thirsty is a reason for Jimmy to drink water in circumstances where there is no serious cost to drinking water that outweighs the utility of quenching his thirst. Nothing in that description seems to be irreducibly normative. Feeling thirsty is obviously not a normative phenomenon, and neither is the act of consuming water, nor are the circumstances I described above, since utility can be described in non-normative terms without remainder.[43] What I take to be irreducibly normative about the reason relation is the relation itself, not the relatum. When something is a reason to do something else, that thing *counts in favor* of that action (Dancy 2006, 39). Counting in favor seems to be the essence of the reason relation that obtains between the variables in the formalism above (Bedke 2009).

What exactly is it for something to count in favor of something else? I am at a loss when I try to analyze the concept into constituent, non-normative parts. As it stands, a favoring relation seems to be a normative relation. To stand or count in favor of something does not seem like any non-normative relations we encounter in the world. Spatial and temporal relations are certainly not like it; neither are familial relations or relations like love. While these relations may count in favor of some action, they are not the favoring relation. Favoring does not seem to resemble mental relations either, like

---

[42] Cf. (Bedke 2009).
[43] Assuming utility is something like the aggregation of certain mental states or interests, none of that strikes me as irreducibly normative.

desiring that P. Desiring that P may be a factor that stands in favor of making P true through your actions, such as a desire that you quench your thirst stands in favor of making it true that you quench your thirst, but the relation between desire and quenching your thirst is not the favoring relation. To see why, consider our Moorean test from section four. It seems strange to say, "I have a reason to do this, but nothing counts in favor of doing it". We can fill in, "this" and "it" with whatever we want, and it will still seem strange. But it does not seem strange if somebody says, "I desire to do this, but there is no reason for me to do it". It seems possible that somebody could desire to do something but say that there is no reason to do it. For example, imagine if somebody has a desire implanted in her mind that makes her want to murder redheads. She realizes that she has this desire and says, "I desire to murder redheads, but there is no reason for me to murder them." I do not see this as strange or reminiscent of the Moorean paradox like the sentences I analyzed in section four. Furthermore, it seems to make sense to say, "I do not desire that P, but there is nonetheless a reason to bring about P." The favoring relation can obtain in the absence of desires, so the relation cannot be reduced to the desire relation.

So there seems to be something irreducibly normative about the favoring relation. Attempts at reducing the relation to something non-normative, like the desiring relation, seem implausible. But Humeans about reasons shouldn't worry about this, because their view can be explained in terms of an irreducible favoring relation. To be a Humean on this account of reasons is to take desires to be the fundamental relatum that takes the place of F in the formalism above. No matter what, F will always be a desire for the Humean. We can thus accommodate views like Humeanism within the

framework of a favoring relation account of reasons. I take this to be a point in favor of this analysis of the reason relation, as it can still make sense of many of the debates between Humeans and anti-Humeans that continue to occur in the literature.

Attempting to reduce the favoring relation to something purely psychological seems not to work so well, as I've discussed above with the example of desire. Motivation may seem like a promising candidate for a reduction base, but it won't work. Rather than being a case of reduction, it seems like motivation can only function as a replacement for the favoring relation. By trying to replace talk of favoring with talk of motivation, we seem to leave something out. Motivation seems to leave out the normativity of reasons that favoring captures quite nicely. Being motivated to do something may be at odds with what we ought to do, as is the case with weakness of the will. We realize that we have reasons to do something that we are either not motivated to do or motivated not to do. But if the favoring relation is just motivation, then we cannot evaluate such situations in line with our common-sense way of talking about motivation and reasons. We encounter situations where motivation conflicts with our normative reasons all the time, and that seems to be evidence that the two are distinct. Just think of the person who is motivated to continue smoking but acknowledges that she has good reasons to stop smoking, or the person who just cannot bring himself to go vegan despite seeing that there are strong moral reasons to stop consuming animal products. These sorts of cases suggest that motivation cannot capture the idea of a reason for action.

Maybe motivation won't work as a reduction base for the favoring relation, but there could be other candidates. Perhaps the favoring relation just is the relation that

obtains between an action and furthering one's ends or goals. On this account, somebody has a reason to do something just if doing it will further her ends or goals. The relation is just that between action and effects that fall in line with something the agent wants, values, or has set for herself as a goal. The problem I see with this view is that it has implausible implications when it comes to disagreements. If somebody has the goal of harming a cat to hear its cries of pain and torturing the cat with a coat hanger heated up on a stove will make it likely that the cat will cry out in pain, then that person has a reason to torture the cat. But what if somebody sees this and says that the person does not have a reason to further her goal of hearing the cat's cries. If reasons are just some sort of statistical relation between action and effect indexed to the goals or ends of an actor, then either the person objecting is conceptually confused, the actor being evaluated is missing some information, or the person objecting lacks the concept of a reason entirely (Bedke 2009). None of these alternatives are palatable as far as I can tell. Claiming that somebody lacks the concept of what you want to reduce because she objects to your reduction seems to be a desperate move, as is claiming that the person must be conceptually confused. As far as the actor lacking information, it does not seem like there are additional non-normative facts that she is unaware of that would influence her goals or methods of achieving them. Such moves also rule out the possibility of genuine and sincere disagreement about whether an actor has a reason to pursue her ends or goals. But such disagreements certainly seem possible and actual, which renders this form of reduction highly problematic (Bedke 2009).

So, it seems like reasons are irreducibly normative relations that can obtain between many kinds of things. The relation is irreducibly normative because it is a

favoring relation, which means that it is a relation where something *counts in favor* of another thing. I went over some possible ways of reducing the favoring relation and concluded that they either outright fail or replace the concept of a normative reason with something purely psychological, which will not work. It seems, then, that moral concepts are essentially connected to irreducibly normative entities, namely reasons. In other words, moral concepts are connected to the concept of a reason for action, which aims to pick out irreducibly normative favoring relations. But as I will show in the next part of this thesis, irreducibly normative relations cannot exist. So, if I am right, then all moral propositions are false since they imply irreducibly normative relations exist when they do not.

Chapter Three

No Irreducible Normativity

1. Preliminaries

In this part of my thesis I will argue that there are no irreducibly normative properties.
Since there are no irreducibly normative properties, and reasons are supposed to be
irreducibly normative properties, then there are no reasons. If moral concepts commit
their users to falsehoods when used to express moral propositions, then moral
discourse is defective. In other words, no moral propositions are true, which is the
formulation of the moral error theory that I argued for in the first part of this thesis. My
strategy will be to defend Bart Streumer's argument against irreducibly normative
properties against objections, some of which he did not address or anticipate in his
recent book. In section two, I will lay out Streumer's argument and explain its
assumptions. In section three, I will address a possible way of avoiding the argument's
conclusion by accepting a principle of property individuation based on a criterion of
meaning. In section four, I will address an objection to the argument based on
hyperintensionality. In section five, I will address a challenge to the argument based on
there not being disjunctive properties. In section six, I will address possible attempts at
rejecting the supervenience of the moral on the descriptive. Finally, in section seven, I
will wrap everything up in this part of the thesis.

2. The Reduction Argument

In this section I will lay out Streumer's reduction argument which has been adapted from
Frank Jackson's work (Streumer 2017; Jackson 1998). The basic idea is that normative

predicates and certain descriptive predicates are necessarily coextensive. Two predicates are necessarily coextensive just in case they have the same extension at every possible world. But if two predicates are necessarily coextensive, then they ascribe the same property. So, normative predicates and certain descriptive predicates ascribe the same properties. And identity of property is surely sufficient (even if not necessary) for reduction. Since normative reasons are supposed to be irreducibly normative properties, they cannot exist, since there are no irreducibly normative properties. So, there are no normative reasons.[44] I will now unpack this condensed version of the argument and explain its assumptions.

First, let's look at normative and descriptive predicates. Consider a particular pleasure state PL1 that has the relational property of being a reason to act a certain way. Anything with normative properties also has descriptive properties, such as having a particular qualitative feel (Streumer 2017, 9). Call those descriptive properties of pleasure, P1-Pn, and the objects O1-On that have those properties Po1-Pon. PL1 thus satisfies the descriptive predicate, D1:

D1(S): S has descriptive properties P1-Pn and for every object O, O is in state S if and

only if O has descriptive properties Po1-Pon.[45]

Since a predicate composed of descriptive predicates is a descriptive predicate, D1 is descriptive (Streumer 2017, 9).

---

[44] Streumer's argument stops at the reduction of normative properties to descriptive properties, but I add that since normative reasons are supposed to be irreducibly normative, they cannot exist given Streumer's argument.

[45] This definition differs from Streumer's definition, which appears to be deficient because in his definition, O1-On do not occur on the left-hand side but they do on the right. Here is Streumer's definition: "**D1**: has descriptive properties P1-Pn, and is such that object O1 has descriptive properties Po1.1-Pon.1, and object O2 has descriptive properties Po1.2-Pon.2 …." (Streumer 2017, 9).

Now consider that pleasure states PL1, PL2, … are all of the reasons for action across all possible worlds. As PL1 satisfies D1, so will PL2 satisfy D2, and PL3 will satisfy D3, etc. All these pleasure states thus satisfy the predicate D*:

D*: satisfies the predicate D1, or D2, or D3, or …

This is also a descriptive predicate since it is composed of descriptive predicates. Now consider this principle:

S: For all possible worlds W and W*, if the instantiation of descriptive properties in W

and W* is the same, then the instantiation of normative properties in W and W* is

the same (Streumer 2017, 10).

This principle will be quite familiar to anybody who has encountered the metaethical literature from the past few decades. Principle S expresses a form of supervenience, which is the idea that some properties depend on others such that if the latter properties are instantiated then so are the former. Now consider this principle:

N: Two predicates ascribe the same property if and only if they are necessarily

coextensive (Streumer 2017, 11).

We can now complete the argument. If **N** is the correct criterion of property individuation, then the predicate, 'is a reason for' and the predicate, 'is pleasurable' ascribe the same property.[46] I used pleasure as a stand-in for any normative property and predicate, meaning that the argument generalizes to all normative properties. Since the argument generalizes, normative reasons can now be our target. Given principle **S**, if the property, 'is a reason for' is instantiated at worlds **W** and **W***, then so is the

---

[46] I am using pleasure as a stand-in for any descriptive property on which normative properties supervene. My argument does not turn on axiological hedonism being true. I will now generalize to all reasons.

property, **D***. For any two worlds, the normative predicate, 'is a reason for' will ascribe a property to an object in every world where the descriptive predicate, 'is **D***' ascribes a property to that object. But, given principle **N**, the properties ascribed by, "is a reason for" and, "is **D***" must be identical, since the two predicates are necessarily coextensive. So, the property, 'is a reason for' is identical to the property, 'is pleasurable'.

The argument is quite simple and easy to understand. It utilizes principles that seem rather plausible, especially supervenience. While I used the descriptive property of pleasure, that was just as an example. The argument works for any descriptive property that we take to instantiate a normative property. So, the argument entails that there are no irreducibly normative properties. In the next section, I will deal with one potential way of avoiding the conclusion of this argument. Then I will address further critiques in the following sections.

3. Individuating Properties by Meanings

One way to avoid the reduction argument is to challenge the principle **N**, which will be a theme among many of the critiques I will address in this part of the thesis. Principle **N** is a criterion of property individuation. The principle is intensional because it individuates based on necessary coextension rather than coextension at the actual world (Audi 2016). However, we could get even more fine grained with how we individuate properties. Instead of necessary coextension we could individuate properties based on the meanings of the predicates that ascribe them. Call this principle M:

M: Two properties are identical if and only if the predicates that ascribe them have the

same meaning.

Now, since we are working within a realist framework when it comes to properties, the properties must not be the meanings of predicates. If the meanings of those predicates just are the properties they ascribe, then we have left the realm of property realism, and therefore we have left the realm of realism about moral properties, which is what the error theory takes to be the correct view when it comes to how we use moral language. We use moral language in a way that commits us to mind-independent moral properties (cf. Joyce 2001). So, principle **M** must express a criterion of property identity that does not take properties to be identical to the meanings of predicates.

Why think that principle **M** is true? It does not seem very plausible when we consider the case of natural kinds. For example, the predicate, 'is water' and the predicate, 'is H2O' seem to have different meanings, but they ascribe the same property. If **M** is true, then those predicates must ascribe different properties (Audi 2016). Besides those sorts of counterexamples, the principle opens up the possibility of a debunking argument. It seems like a massive coincidence that the meanings of our predicates would perfectly match the structure of the world. Why think that the meanings of our predicates always match up with mind independent properties? If we take the meanings of our predicates to be determined by the world in the way a semantic externalist would think, then principle **M** won't work as a criterion of property identity. If we must look to the world to see which properties are regulating the meanings of our predicates, then we cannot look to our predicates' meanings to determine which properties are identical. So, principle **M** must function alongside a view

that takes meanings to be in the head in some sense. But once we locate meanings as an internalist does, we must address the massive coincidence that **M** represents.

The proponent of principle **M** has a burden of proof to explain how the meanings of our predicates perfectly match up with mind-independent properties. Again, since we are assuming anti-realism is false, the properties individuated by **M** are independent of the meanings of the predicates that ascribe them, yet those predicates perfectly match up with those properties. Why is that? It may be possible to explain the coincidence when it comes to mental properties that we are intimately familiar with; they may have played a role in the formation of those predicates and their meanings.[47] It may be plausible to argue that our direct acquaintance with certain qualitative mental states played a role in the formation of our predicates about them. Qualitative pain states probably played a role in the formation of the predicate, "is painful" by virtue of our direct acquaintance with it. However, when it comes to other sorts of properties, that sort of story becomes farfetched. Properties we may never encounter yet we can meaningfully speak of, such as being a three-mile-wide piece of gold do not feature in the formation of our predicates about them, yet given principle **M**, our predicates perfectly track them. Such a phenomenon needs an explanation, and the proponent of **M** must provide it.

One possible way of defending principle **M** is to argue that we can make sense of the perfect symmetry between predicate meanings and properties by introducing functions from simple to more complex predicates. If **F** and **G** are meaningful predicates that stand for properties **P** and **Q**, respectively, then, "**F & G**" is a complex predicate that expresses the property **CONJ(P, Q)**, where **CONJ** takes pairs of properties to

---

[47] See (Chalmers 2006) for more on this possibility.

conjunctions of those properties.[48] The problem with this move is that the simple properties picked out by simple predicates will have to exhaust the property-base out of which more complex properties are composed.

The picture painted by this move would be plausible if there was a finite number of simple properties that our predicates' meanings about them perfectly track, and all of our complex predicates are built out of those simple predicates. But it still seems like such a picture must explain why our predicates' meanings perfectly match all of the simple properties out of which the more complex properties are constituted. Even if we can explain how a complex property like being a three-mile-wide piece of gold is constituted by the properties of being three miles wide and being a piece of gold, and then show how those simpler properties match up with our simple predicates plus a conjunction relation, we are stuck with a coincidence in need of explanation. Why is it that our simple predicates perfectly match all of the simple properties in the world? Until the proponent of **M** can explain how the meanings of our predicates perfectly match the simple properties distributed across possible worlds, her view will seem quite mysterious.

To conclude this section, the proponent of principle **M** must deal with counterexamples like water and H2O being identical while their respective predicates lack synonymy. She must also explain how the meanings of our predicates perfectly match up with the properties that are distributed across possible worlds. While there may be a way to show how, given a base of simple properties and predicates, we can build up more complex properties and predicates, such a view is still stuck with an

---

[48] Thanks to Graham Oddie for pointing out and articulating this possibility.

explanatory burden when it comes to how the meanings of our simple predicates perfectly mirror the distribution of simple properties.

4. The Challenge of Hyperintensionality

Principle **N** is supposed to express a necessary truth about properties, namely that necessarily coextensive predicates stand for the same property. If we have an example where two predicates are necessarily coextensive, yet they stand for different properties, then principle **N** is false. Consider this counterexample: the predicates, "is Socrates" and, "is the member of the singleton set {Socrates}" are necessarily coextensive, yet it seems like the property of being Socrates and the property of being the member of the singleton set, {Socrates} are distinct. Examples like this are instances of hyperintensional property individuation (Audi 2016). If the Socrates case is a genuine counterexample to principle **N**, then **N** fails to account for the fine-grained structure of the world. Such structure is supposed to be composed of non-causal determination relations among properties.[49] When one property non-causally determines another, the prior property is said to be more fundamental than the latter property, and therefore it provides a metaphysical explanation of the asymmetry between the two properties (Audi 2012). The non-causal determination relation is meant to explain why there is an asymmetry between properties, where one determines the other and not vice versa. So, the property of being Socrates is supposed to explain the property of being the member of the singleton set, {Socrates} in terms of the former grounding the latter (Audi 2012).

---

[49] Other candidates besides properties include facts, but I will continue using property-talk.

The intuition that guides accounts of grounding is the asymmetry between the allegedly distinct properties like being Socrates and being the member of, {Socrates}. The apparent asymmetry is therefore the datum that stands in need of an explanation, which some theorists think can be accounted for with the grounding relation. While I feel the force of the intuition about there being some sort of asymmetry between the *predicates*, "is Socrates" and, "is the member of the singleton set, {Socrates}", I do not think that we need a grounding relation to explain this phenomenon. I think that there is another explanation of the asymmetry available to us that does not require us to posit the grounding relation as a mind-independent phenomenon.[50]

The example of Socrates and {Socrates} can be interpreted in a way that does not require us to posit a worldly grounding relation. I do not think that the predicates, "is Socrates" and, "is the member of the singleton set, {Socrates}" ascribe different properties. It seems like the same property is being picked out by those predicates; to be the member of the singleton set, {Socrates} and to be Socrates is the same property. The predicates are just different ways of picking out the same property, one more complicated than the other. What makes the one more complicated than the other is itself a complicated issue, but whatever it is, it will be something epistemic or semantic rather than ontic.[51] I take it that what makes the predicates complicated will differ based on the case we are examining. In the Socrates, {Socrates} example, one possibility is that there are more words in the predicate about set membership which do not do any actual work in picking out worldly features over and above the predicate about

---

[50] For anti-realist accounts of grounding, see (Thompson MS).
[51] I think the best and most simple answer is that they involve different words and/or concepts. I take concepts to fall within the field of semantics, but that's an unimportant terminological issue.

Socrates.[52] In this case, then, the different words generate the asymmetry between the two predicates and lend credence to the illusion of there being two different properties rather than two different ways of zeroing in on the same property.

Another kind of case involves the predicates focusing our attention onto different parts of the same property. Take the example of trilaterality and triangularity. The predicates that ascribe those properties are supposed to be necessarily coextensive, yet the properties seem to be distinct (Audi 2016).[53] First off, I do not think that the predicates that ascribe those properties are necessarily coextensive. There can be figures with three sides that are not triangles, such as an open shape with three sides whose two interior angles amount to one hundred eighty degrees (Streumer 2017, 13).[54] So, trilaterality and triangularity are distinct properties. Let's fix the example up a bit. How about the predicates, "is a closed, three-sided figure" and, "is a closed, three angled figure"? These predicates seem to be necessarily coextensive, but do they ascribe the same property? I think they do. If we were to draw the figure picked out by each predicate, we would get triangles in both cases. All that the predicates do is draw our attention to different aspects of the figure, namely its three sides and its three interior angles.[55] So the trilaterality/triangularity example fails to refute principle N as well.

Maybe we can find a hyperintensional counterexample elsewhere. Here is an example: the number two and the square root of four (Parfit 2011, 296-297). I fail to feel the force of this alleged counterexample. The predicates that ascribe those properties

[52] Thanks to Bart Streumer for talking this issue through with me.
[53] The figures discussed are in Euclidean space.
[54] Picture a square without the top line on it.
[55] N.b. There doesn't seem to be any asymmetry here, either. I can't tell which is supposed to ground which, the three sidedness or the three angledness of the closed figure.

both refer to the number two. All that the predicate, "is the square root of four" does is draw our attention to a way of arriving at the number two through a mathematical procedure; it does not pick out a distinct property from being the number two. Ok, maybe a different example: consider the predicates, "is a round square" and, "is an even prime number larger than two" (Olson 2014, 93). Both of these predicates are necessarily coextensive, since they necessarily lack extension. Perhaps, then, they ascribe different properties to objects yet are necessarily coextensive. I do not think that it makes sense to say that predicates that necessarily lack extensions ascribe properties to anything. If such predicates ascribe anything, they will be non-existent properties. But the question of whether two objects have the same non-existent property seems nonsensical, so this counterexample cannot work. Properties are ways that objects can be, and to say that there are non-existent ways that objects can be makes no sense (Streumer 2017, 17-18).

5. A Challenge to predicate D*

Perhaps predicate **D*** fails to stand for a genuine property because it is an infinitely long disjunction. The fact that it is, at least potentially, infinitely long is not really a problem. The objector may argue that since we cannot translate the predicate into a natural language like English, it fails to be a genuine predicate. But that can't be right; just because the predicate cannot be translated into a natural language does not mean that it is not a genuine predicate. To argue that way is to exhibit linguistic chauvinism towards natural languages. Since there is a possible language in which predicate **D***

can be expressed, it seems fine to say that it is a genuine predicate (Streumer 2017, 26).

Another challenge to predicate **D*** is based on the view that there are no disjunctive properties. Predicate **D*** does not ascribe a property because there are no disjunctive properties. But properties are not disjunctive, only the predicates that ascribe them are or can be. Given principle **N**, there are properties that can be ascribed both disjunctively and non-disjunctively, like, "is a noble gas" (Streumer 2017, 27). Such examples show that disjunctive predicates can be replaced with non-disjunctive predicates without changing their extensions, which is evidence that disjunctivity is a phenomenon that has to do with the nature of our predicates and not worldly properties.

Another way of arguing that **D*** does not pick out a genuine property is developed by Graham Oddie (Oddie 2005, 153-155). According to Oddie's view, a predicate picks out a genuine property only if it carves out a convex region of quality space. A region **R** is convex if and only if the region between two sub-regions of **R** is itself **R**. So, for example, being either hot or cold is not a genuine property because there is a sub-region between them that is neither hot nor cold, namely warm. Maybe predicate **D*** fails to carve out a convex region of quality space, so it fails to pick out a genuine property. I do not think that this move succeeds for two reasons. First, I am not entirely sure what quality space is supposed to be. According to Peter Gärdenfors, whose view of property individuation Oddie works from, the space in which predicates carve out convex regions is called conceptual space, which seems like a case of properties being concepts, which is not a form of realism about properties (Gärdenfors 2000 & 2014).[56]

---

[56] In the sense required for those properties to be real moral properties according to popular versions of moral realism. No moral realist I am aware of would say that moral properties are identical to concepts.

If quality space is not the same thing as conceptual space, then it must be worldly, which brings us to my second problem.[57] It seems like moral predicates like, "is right" could be necessarily coextensive with descriptive predicates if certain first order moral theories are true. So, for example, if a simple version of hedonistic act utilitarianism is true, then moral predicates like, "is right" would be necessarily coextensive with the descriptive predicate, "maximizes total pleasure" and therefore, according to principle **N**, they ascribe the same property. But it seems false that the issue of reductionism and non-reductionism in metaethics should turn on which first order moral theory turns out to be true. It seems to me that the question of whether moral properties are irreducible or reducible to descriptive properties turns on the nature of those properties rather than which first order moral theory turns out to be true. Furthermore, it seems like proponents of first order moral theories need not adopt a particular metaphysical picture of the properties they are discussing, so the two issues are independent of each other. This leads me to doubt that Oddie's account is not an appropriate way to address the question of moral property individuation (Streumer 2017, 28).

6. Challenges to Supervenience

So hyperintensionality failed to give us good counterexamples to principle **N**, and predicate **D*** stands for a genuine property. Maybe another way out of the reduction argument is to challenge the principle **S**, supervenience. Recall what principle S is:

---

[57] A third problem I do not want to pursue here is that Oddie's view seems to entail that all properties come in degrees, which excludes categorical properties. I will not push this issue here since it complicates everything too much.

S: For all possible worlds W and W*, if the instantiation of descriptive properties in W and W* is the same, then the instantiation of normative properties in W and W* is the same (Streumer 2017, 10).

Principle **S** seems extremely plausible, to the point where some theorists think that accepting it is a mark of being a competent moral concept user (Streumer 2017, 6-7). One reason to reject supervenience is because it does not allow for there to be fundamental moral properties or facts. Principle S seems to imply that moral properties depend on descriptive properties, which means that moral properties will always be less fundamental than at least some descriptive properties.[58] But if there are moral properties at the most fundamental level of the world, then principle S must be false.[59] The idea that there is a fundamental level of reality that is moral strikes me as very implausible. If there are such fundamental moral facts or properties, then there could be two descriptively identical worlds where a particular action is wrong in the first world and permissible or right in the second world. In other words, we will have brute moral facts that cannot be explained in terms of the presence or absence of some descriptive property or other. We would end up with worlds where every actor ought to act such that she maximizes the average goodness in that world, and other worlds where every actor ought to act such that she maximizes the total goodness in that world, where both of those worlds are descriptively identical. All that would differ between those worlds are the fundamental moral facts.

---

[58] Fundamentality is both relative and comes in degrees (Audi 2012).

[59] See (Oddie 2005, 146-166) for a view that aims at preserving both **S** and **N** while denying reduction. While Oddie's argument against reductionism does not require us to adopt the view I address in this section, his developed view is relatively similar to the one that I address here. Oddie's view rejects Boolean closure, meaning that a condition is specifiable in terms of purely natural properties but may not be a purely natural property. The Boolean operators are not property preserving, but rather what is property preserving is the ability to carve out a convex region of what Oddie calls, "quality space". I addressed Oddie's view of property individuation in the previous section.

The view I described above strikes me as extremely implausible, but I do not have a knockdown argument against it. I think that it is much less plausible than principle **S**, which means that the opponent of **S** shoulders the burden of proof here. She must give us good reasons to think that there is a fundamental moral level of reality, since that claim implies that an extremely widespread principle in metaethics, accepted by anti-realists and realists alike, is false (Miller 2014, 29). I will leave this issue here by noting that if this is the cost of abandoning principle **S**, then I would choose supervenience.

Another way of challenging supervenience is to provide an explanation for why it seems to be the case that principle **S** is true, where that explanation does not involve the truth of supervenience. One way to do this is to consider the alleged supervenience relation as analogous to the causal relation. If we take an orthodox Humean approach to supervenience as we can with causation, we might be able to explain why it seems to be that principle **S** is true without appealing to necessary connections between moral and descriptive properties (Hills 2009). The orthodox Humean considers the causal relation to be nothing over and above a relation of constant conjunction (Morris and Brown 2014). Analogously, we may consider what we take to be supervenience to actually be a relation of constant conjunction between certain descriptive properties and moral properties. For example, we may repeatedly encounter pleasure in the world which is constantly conjoined with the property of goodness, and then we may erroneously infer that goodness supervenes on pleasure, when in reality the two just happen to be coinstantiated in this world. What this view entails is that there may be worlds that are exactly like ours descriptively, yet they radically differ morally. So there

may be a world just like ours, except pleasure is bad and pain is good. There is no necessary connection of determination between pleasure and goodness or pain and badness on this view, so there is no modal restriction on the instantiation conditions for moral properties. Such a view entails that principle **S** is false, since **S** posits a necessary connection where there is none.

I have a similar problem with this view as I do with the previous account that denies supervenience. The problem is that this view entails that it is not only a brute fact that pleasure is good, but there may be worlds exactly like ours except pleasure is bad. Think about it, a world that perfectly resembles ours at the descriptive level, except pleasure is bad and pain is good. This view is extremely farfetched and implausible. Not only that, though, but it seems to restrict our ability to use modally infused moral reasoning. By that I mean a sort of moral reasoning that appeals to various sorts of possibilities that do not obtain in our world. If there are no modal restrictions on the distribution of moral properties across possible worlds in the way that principle **S** entails, then why think that our modally infused moral reasoning is a reliable guide to moral knowledge? It seems like you cannot trust your moral reasoning on this account, as long as such reasoning includes claims about possibilities that we do not encounter in this world.

7. Wrapping Up

To conclude, we have laid out the reduction argument, which, on very plausible premises, entails that there are no irreducibly moral properties. Since there are no irreducibly moral properties, then moral propositions that commit their users to such

properties are never true. I argued in the previous part of this thesis that such moral propositions are pervasive in moral discourse, enough so that it is safe to say that all moral propositions are never true since they all imply that there are irreducibly normative reasons for action. In section two of this part of the thesis, I addressed a potential challenge to the principle of property individuation deployed in the reduction argument based on synonymy as a criterion of property identity. In section three, I addressed the challenge that arises when we consider hyperintensionality as a possible way to individuate properties. In section four, I dealt with two ways of avoiding the principle of the supervenience of the moral on the descriptive. We have now reached the conclusion of this thesis, namely that the moral error theory is true. What to do in light of this finding is a separate question that will be addressed in my future work.

Bibliography

Audi, Paul. "Grounding: Toward a Theory of the In-Virtue-Of Relation." *Journal of Philosophy*, vol. 109, no. 12, 2012, pp. 685–711.

Audi, Paul. "Property Identity." *Philosophy Compass*, vol. 11, no. 12, 2016, pp. 829–840.

Bedke, Matthew S. "Might All Normativity Be Queer?" *Australasian Journal of Philosophy*, vol. 88, no. 1, 2009, pp. 41–58.

Chalmers, David J. "Perception and the Fall from Eden." *Perceptual Experience*, 2006, pp. 49–125.

Cohen, Stewart. "Justification and Truth." *Philosophical Studies*, vol. 46, no. 3, 1984, pp. 279–295.

Edwards, Gary. The Many Faces of Besire Theory. M.A. Thesis. 2011.

Flanagan, Owen J. *The Geography of Morals: Varieties of Moral Possibility*. Oxford University Press, 2017.

Frankfurt, Harry. "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy*, 66: 829–39. 1969.

Fumerton, Richard A. "The Paradox of Analysis." *Philosophy and Phenomenological Research*, vol. 43, no. 4, 1983, p. 477.

Gärdenfors, Peter. *Conceptual Spaces: The Geometry of Thought*. MIT Press, 2000.

Gärdenfors, Peter. *Geometry of Meaning: Semantics Based on Conceptual Spaces*. The MIT Press, 2014.

Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis*, vol. 23, 1963, p. 121-123.

Glasgow, Joshua. *A Theory of Race*. Routledge, 2009.

Grice, H. P. Meaning, *Philosophical Review*, 1957. 66: 377-88.

Hills, Alison. "Supervenience and Moral Realism", in Hieke Alexander & Leitgeb Hannes (eds.), *Reduction, Abstraction, Analysis.* 2009.

Horgan, Terry, and Mark Timmons. "What Do Reasons Do?." *Metaethics after Moore*, 2006, pp. 39–60.

Horn, L. R. Implicature, in L. R. Horn & G. Ward (ed.), *The Handbook of Pragmatics*, 2004, pp. 3–28. Oxford: Blackwell Publishing.

Jackson, Frank. *From Metaphysics to Ethics: a Defence of Conceptual Analysis*. Clarendon, 1998.

Joyce, Richard. *The Myth of Morality*. Cambridge University Press, 2001.

Kaspar, David. "Intuitionism and Nihilism." *Philosophia*, 2017, pp. 1-18.

Loeb, Don. Moral Incoherentism: How to Pull a Metaphysical Rabbit out of a Semantic Hat." Walter Sinnott-Armstrong, ed., Moral Psychology, Vol.2, Cambridge, MA: MIT Press, 2008

Mackie, John Leslie. *Ethics: Inventing Right and Wrong*. Penguin Books, 1977.

Marks, Joel. *Ethics without Morals: in Defence of Amorality*. Routledge, 2013.

McAuliffe, William. Do Emotions Play a Central Role in Moral Judgments, MS.

Miller, Alexander. *An Introduction to Contemporary Metaethics*. Polity, 2003.

Moore, George Edward. *Principia Ethica*. Cambridge University Press, 1903.

Morris, William Edward and Brown, Charlotte R., "David Hume", *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/hume/>.

Oddie, Graham, "Truthlikeness", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/truthlikeness/>.

Oddie, Graham. *Value, Reality, and Desire*. Clarendon Press, 2005.

Olson, Jonas. "In Defense of Moral Error Theory." *New Waves in Metaethics*, 2010, pp. 62–84.

Olson, Jonas. *Moral Error Theory: History, Critique, Defence*. Oxford University Press, 2014.

Parfit, Derek, and Samuel Scheffler. *On What Matters*. Vol 2. Oxford University Press, 2011.

Parsons, J. "Cognitivism about Imperatives." *Analysis*, vol. 72, no. 1, July 2011, pp. 49–54.

Pierce, Jeremy. *A Realist Metaphysics of Race: A Context-Sensitive, Short-Term Retentionist, Long-Term Revisionist Approach*. Lexington Books, 2015.

Pigden, Charles R. "Nihilism, Nietzsche, and the Doppelganger Problem." *A World Without Values*, Sept. 2007, pp. 17–34.

Putnam, Hilary, 1975. "The Meaning of Meaning," *Philosophical Papers, Vol. 2: Mind, Language, and Reality,* Cambridge University Press.

Sayre-Mccord, Geoffrey. "The Many Moral Realisms." *The Southern Journal of Philosophy*, vol. 24, no. 1, 1988, pp. 1–22.

Schroeder, Mark. *Noncognitivism in Ethics*. Routledge, 2010.

Sinnott-Armstrong, Walter. *Moral Skepticisms*. Oxford Univ. Press, 2006.

Smyth, Nicholas. "The Function of Morality." *Philosophical Studies*, vol. 174, no. 5, 2017, pp. 1127–1144.

Sperber, Dan and Deirdre Wilson. *Relevance: Communication and Cognition 2nd ed.* Blackwell Publishing, 1995.

Streumer, Bart. *Unbelievable Errors: An Error Theory about All Normative Judgements*. Oxford University Press, 2017.

Thompson, Nancy, "Irrealism about Grounding", MS.

Tittle, Peg. *What If ---: Collected Thought Experiments in Philosophy.* Pearson Longman, 2004