

Moral Properties and Moral Imagination

by

Tyler Doggett

A.B. Philosophy
Princeton University, 1998

SUBMITTED TO THE DEPARTMENT OF LINGUISTICS AND PHILOSOPHY IN
PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN PHILOSOPHY
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

SEPTEMBER 2004

© 2004 Tyler Doggett. All rights reserved.

The author hereby grants to MIT permission to reproduce
and to distribute publicly paper and electronic
copies of this thesis document in whole or in part.

Signature of Author: _____


Department of Linguistics and Philosophy
July 23, 2004

Certified by: _____

Alex Byrne
Associate Professor of Philosophy
Thesis Supervisor

Accepted by: _____

Vann McGee
Professor of Philosophy
Chair, Committee on Graduate Studies



ARCHIVES

SECRET
1950 OCT 20

SECRET
1950 OCT 20

Moral Properties and Moral Imagination

By

Tyler Doggett

Submitted to the Department of Linguistics and Philosophy
on July 23, 2004 in Partial Fulfillment of the Requirements for
the Degree of Doctor of Philosophy in Philosophy

ABSTRACT

“Moral Realism” is about an argument against moral realism, why it is unsound, and what emerges from that. The argument is that if there were moral properties, they would be queerly related to non-moral properties and this is sufficient reason to think there are no moral properties. The argument is unsound for two reasons. The first emerges from consideration of sensational properties like being in pain or being in ecstasy—they bear the queer relation to non-sensational properties. The second emerges from consideration of vice properties like being an instance of greediness—they are not queerly related to non-moral properties. Analogies between moral and sensational properties are discussed.

A disanalogy between the moral and sensational is important to “The Explanatory Gap” which discusses Levine’s notion of an explanatory gap, relates it to the queer relation discussed in “Moral Realism,” and criticizes one use to which it is put. The criticism emerges from consideration of the disanalogy between the moral and sensational: our moral imagination is considerably more limited than our sensational imagination.

That there are limits to our moral imagination is interesting. “Imaginative Resistance” solves an old puzzle from Hume about the limits of our imaginative capacities, for example, the inability of some people (myself, for example) to imagine that baseless killing is morally permissible. Both the puzzle and solution illuminate the natures of imagination and possibility and the relation between them.

Thesis Supervisor: Alex Byrne
Title: Associate Professor of Philosophy

ACKNOWLEDGMENTS

I received a lot more help with this dissertation than I deserved. (I do not know exactly how much I did deserve.) Many people heard or read drafts of it and gave me very helpful comments. Thanks to them all.

I am most grateful to the members of my committee—Alex Byrne, Judith Jarvis Thomson, and Stephen Yablo—and to Selim Berker, Andy Egan, Ned Hall, James John, Alison Simmons, and Robert Stalnaker. I have taken so many examples, ideas, and phrases from them that I have decided not to credit them in the text—it would embarrass me to see how little contribution I have made to these papers.

Thanks to my fellow graduate students, the department staff, and my teachers for making MIT such a calm, encouraging, stimulating, and wondrous place to do philosophy.

Above all, thanks to the linguistics half of the department for admitting Teal Bissell who made our home such a calm, encouraging, stimulating and wondrous place to do philosophy.

CHAPTER ONE MORAL REALISM

Let moral realism (for short, realism) be the following conjunctive thesis:

MORAL REALISM—There are mind-independent moral properties and they are actually instantiated.

This is a minimal view and different realists add different theses to it. Some add that moral properties are non-natural properties. Some add, by contrast, that moral properties are natural properties and also that they have causal powers. Some agree that they are natural but deny that they have causal powers. But all realists accept that there are mind-independent moral properties, actually instantiated.

Relatedly, all anti-realists deny that moral properties are mind-independent or deny that there are any moral properties or deny that any are instantiated. Subjectivists do the first. They say that there is a moral property goodness, actually instantiated but it is mind-dependent: it is the property of being desired or being desired to be desired. Relativists, too, say that goodness is mind-dependent: it is the property of being believed to be good by members of a community or being held to be good by a bargain the members made. Realists deny that goodness is like any of those properties. The realist G.E. Moore, for example, thinks that goodness can be instantiated in a world with no minds at all. From now on, I mean ‘mind-independent moral properties’ by ‘moral properties.’

Emotivists disagree with realists, relativists, and subjectivists alike. They deny that there are any moral properties, mind-independent or otherwise. They say that there is no rightness, for example, or wrongness or any other moral property.

Finally, an anti-realist can disagree with emotivists, realists, relativists, and subjectivists. He can accept that there are moral properties, so long as he denies that any are instantiated. He can say that there is rightness; it is instantiated in some possible worlds but not in the actual world.

All anti-realists agree, though, that it is not the case that moral properties are actually instantiated. There is a powerful argument that they are right; the argument follows.

1. THE ARGUMENT AGAINST REALISM

The argument starts with

- (1) If any moral property is instantiated, its instantiations are determined by the instantiation of natural properties.

The instantiation of a property *x* is determined by the instantiation of a property *y* just in case the instantiation of *y* suffices for the instantiation of *x* and this is true across all possible worlds, not just the actual world. If things instantiate moral properties then, in the sense just given, their instantiation is determined by the instantiation of the thing's natural properties.

'Natural properties'? We can include, at least, all the non-moral properties. Maybe the moral properties themselves are natural but, as noted in the introduction, that is controversial. Being an instance of burning a cat just for fun, for example, is a natural property. Say that Alan burns a cat just for fun (for short, Alan burns a cat). If any action is bad, burning a cat is bad (not just bad for the cat or bad for Alan's conscience but bad, period). If Alan's action instantiates badness—a moral property—this is determined by its instantiating being an instance of cat burning—a natural property. This

first premise is not controversial. Realists and anti-realists accept it; they just disagree about whether the antecedent is true.

The next premise is

(2) For any natural property, it is an open question *why* things instantiating that property instantiate a moral property, if they do.¹

It is an open question *why*, for example, when someone burns a cat, his act is bad. It is not an open question *whether* his act is bad. Everyone knows that it is. What is open is *why* this bad act is a bad act.

The notion of an open question comes from G.E. Moore although Moore says very little about what an open question is. Some examples will fix ideas about it. For instance, take the question, “Why is cat burning bad?” Take, also, a conceivable answer to it: cat burning is bad because it hurts the cat. Someone might accept the answer—I do—without thinking it is a perfectly satisfying explanation. The answer just produces a new question, “Why is hurting cats bad?” Take a conceivable answer to that question: hurting cats, and, thus, burning cats, is bad because cats don’t like it. Someone might accept *that* answer while still thinking it too is a less than perfectly satisfying explanation. He might wonder, “Why is doing something that cats don’t like bad?” A trail of why-questions has been opened up and it is hard to see where (or if) it stops. Each answer to each question in the series seems to raise a question analogous to the first. Whatever a perfectly satisfying explanation turns out to be, I assume that it is a necessary condition on being one that it blocks the trail of ‘why’s.

To say that it is an open question why cat burning is bad is to say that there is no such perfectly satisfying explanation of why it is bad. The claim that it is an open

¹ If moral properties are natural properties, then this premise should begin: ‘for any non-moral natural property.’

question why cat burning is bad is not the claim that it seems that cat burning is not bad. It is also not that it seems arbitrary that cat burning is bad. It is also not that any answer to “Why is cat burning bad?” seems wrong or arbitrary. It is just that for any conceivable answer to the question, it seems in order to ask why it is true. If we had a good explanation for why cat burning is bad, this would not be so. This becomes clear from considering cases where there is no open question.

Take a square with one-foot long sides. It is not an open question why something instantiating the property of having four one-foot long sides and four right angles instantiates the property of being a square. That it does is guaranteed by the meaning of ‘square.’ The trail of why-questions is blocked once you know the meaning of the word. So someone who says, “I wonder why squares have four sides” does not know the meaning of ‘square.’ Someone who wonders why cat burning is bad might perfectly well know what ‘bad’ (and ‘cat burning’) mean. When you teach someone what ‘bad’ means (or ‘cat burning’), you do not thereby teach them why cat burning is bad. When you teach them what ‘square’ means, you do teach them why the thing with four one-foot long sides and four right angles is a square.

By further contrast with the moral case, consider a small glass full of water. It is not an open question why things instantiating the property of being a small glass full of H_2O instantiate the property of being full of water. Unlike “Why are things with four one foot sides and four right angles squares?,” though, “Why is a small glass of water a glass of H_2O ?” is not answerable by just anyone who understands the meaning of its words (I cannot answer it, for example). You can only answer it by learning some chemistry. Learning the requisite chemistry blocks the trail of why-questions. Someone who

wonders why a small glass of water is a glass of H₂O is ignorant of chemistry. Someone who wonders why cat burning is bad is ignorant of what, exactly? Ethics tells you that cat burning is bad, but for any reason it gives as to why it is bad—it hurts, cats do not like it—you can ask why it is true. After all, Moore might well accept that there is an open question about why cat burning is bad even though *Principia Ethica* contains a lot of ethical theory.

The water case brings out a respect in which my notion of an open question differs from Moore. Moore holds that there is an open question about whether an act that is desired to be desired is a good act and this is shown by the fact that *a priori* it makes perfect sense to wonder whether something that is desired to be desired is good. By contrast, Moore argues, *a priori* it makes no sense to wonder whether something good is good. I drop the restriction to what makes *a priori* sense. It is not just that you can wonder *a priori* why cat burning is bad. You can wonder *a priori* why a small glass of H₂O is a glass of water. But, unlike in the water case, you do not stop wondering why cat burning is bad after learning lots of ethics or chemistry or whatever.

There is, then, a case for the thesis that it is an open question why acts instantiating being an instance of cat burning are bad. Some realists will deny the thesis,² but I think there is something to it. At least, there is a contrast between “Why is cat burning bad?” and both “Why is something with four one-foot long sides and four right angles a square?” and “Why are small glasses of H₂O glasses of water?” We have not argued for the thesis that for *any* natural property, it is an open question why things instantiating it also instantiate *any* moral property (if they do). However, in arguing for the thesis that it is an open question why cat burnings are bad, we made no use of features

² It is rejected in Brink (1989), Jackson and Pettit (1995), Jackson (1998), and Sturgeon (2003).

peculiar to being an instance of a cat burning but not other natural properties. We also made no use of features peculiar to badness but not other moral properties. So it is plausible that we may infer

(2) For any natural property, it is an open question *why* things instantiating that property instantiate a moral property, if they do.

We know from (1) that instantiations of moral properties are determined by instantiations of natural properties. Badness is determined, for example, by being an instance of cat burning. So it is an open question why an act instantiating being an instance of cat burning instantiates badness even though the instantiation of the first determines the instantiation of second. For short, the determination here is *brute*.

BRUTE DETERMINATION—For any property *x* and any property *y*, *x* is brutally determined by *y* iff the instantiation of *x* is determined by the instantiation of *y* and it is an open question why things instantiating *y* also instantiate *x*.

From (1) and (2), then, it follows that

(3) If any moral property is instantiated, it is brutally determined by a natural property.

For example, if cat burning is bad, its badness is brutally determined by being an instance of cat burning.

Brute determination is pretty queer. To see in what way queer, consider again the square with one-foot sides. This figure instantiates being a square. The instantiation here of being a square is determined by the instantiation of having four one-foot long sides and four right angles. It is not an open question why things instantiating the second instantiate the first.

Consider, again, the small glass of water. It instantiates the property of being a glass full of water. The instantiation here of being a glass full of water is determined by

the instantiation of being a small glass full of H₂O. It is not an open question why things instantiating the second instantiate the first.

Maybe there is just no explanation for why cat burning is bad; maybe the fact that cat burning is bad is not the kind of fact that can be explained. Maybe—but what are the reasons for thinking it is like that and, if it is like that, why? Why isn't it like the fact that the figure with one-foot sides is a square or that the small glass of H₂O is a glass of water?

What the square and water examples suggest is the general, plausible principle that if the instantiation of one property determines the instantiation of another, then the determination is not brute. There is some explanation for why the instantiation of one determines the other and this explanation blocks the trail of why-questions. That is,

(4) There is no brute determination.

From (4) and

(3) If any moral property is instantiated, it is brutally determined by a natural property,

it follows that

(5) Moral properties are not actually instantiated.

This completes the powerful argument for (5). If the argument is sound, then, since realism is the thesis that there are moral properties, actually instantiated, realism is false.³

³ (1)-(5) is quite a bit like arguments in Blackburn (1973), Blackburn (1984), Blackburn (1985), Hare (1984), Horgan (1993), Horgan and Timmons (1992), Mackie (1977), Mackie (1980), and Mackie (1982). How like those arguments (and how alike they are to one another), I am unsure. One clear difference is that all of the arguments besides (1)-(5) are put in terms of supervenience not determination. The idea behind (1)-(5) and the rest is roughly the same, though. Mackie, for example, complains that if realism is true, then an action is bad *because* of its being a cat burning and it is hard to see what the force of the 'because' is. I take him to be complaining that the determination is brute. He is not complaining about supervenience. Supervenience relates groups of properties—e.g. the moral and the natural—not particular properties like being an instance of cat burning.

2. WHAT IS WRONG WITH (1)-(5)

(1)-(5) is not sound because (4) is not true. If it were, no one would ever be in pain. It will take some work to bring this out.

Let sensational properties include being in pain, being in ecstasy, being in agony, and so on. Most people agree that the instantiation of sensational properties is determined by the instantiation of neural properties or functional properties or behavioral properties or... What is disputed is which do the determining. Identity theorists say, "neural properties," functionalists, "functional properties," and so on. Let us assume that the identity theorists are right and let us assume, further, that the instantiation of being in pain is determined by the instantiation of having your C-fibers fire.

There is a lot to being in pain, including feeling a certain way, being disposed to act a certain way, and being disposed to want certain things. In this way, being in pain is like being a glass of water. There is a lot to being a glass of water, including being a glass of liquid, being a glass of clear stuff, and being a glass of potable stuff. Being a glass full of water is *not* brutally determined by being a small glass full of H₂O because there is no open question about why small glasses full of H₂O are glasses full of water. That is, there is no open question about why they are glasses full of liquid and full of clear, potable stuff. With being in pain, things are different. We can assume that there is nothing brute about why having your C-fibers fire suffices for being disposed to act a certain way and being disposed to want certain things but that is not all that needs

The description of brute determination as queer alludes to the "The Argument From Queerness" section of Mackie (1977). It is a badly named section since, as Mackie admits, there are at least three arguments from queerness made in it, one of which is quite like (1)-(5). A lot more on Mackie is in my manuscript "The Argument From Queerness."

explaining. Why is it that having your C-fibers fire suffices for having any feeling at all? This is an open question. For any conceivable answer, an analogous question arises. As Joseph Levine writes,

There seems to be nothing about C-fiber firing which makes it naturally 'fit' the phenomenal properties of pain...[T]he connection between [being in pain and being a C-fiber firing] is completely mysterious. (Levine (1983): 356)

Note that what is mysterious here is not *whether* the instantiation of having your C-fibers fire determines the instantiation of being in pain. We are assuming that it does. What is mysterious is *why*. Levine writes, more recently,

In the case of a psychophysical [determination] claim it still seems quite intelligible to wonder how it could be true, or what explains it, even after the relevant physical...facts are filled in. (Levine (2001): 82; I have replaced 'identity' with 'determination')

Again, what is intelligible is not wondering *whether*, for example, the instantiation of being in pain is determined by the instantiation of having your C-fibers fire. You might know that it is. What is intelligible is wondering why it is so determined. For any conceivable explanation, it makes sense to wonder why it is true. Levine's view is that being in pain is brutally determined.⁴

The case for thinking it is brutally determined is at least as strong as the case for thinking that, say, badness is. In one way, the case is even better. For example, in the moral case, where being a cat burning brutally determines badness, it does not seem possible that an action instantiates the first without instantiating the second. It seems that any action instantiating the first would instantiate the second. By contrast, it seems possible that your C-fibers are firing but you are feeling nothing at all. This is a mere

⁴ His terminology is that there is an *explanatory gap* between the fact that your C-fibers are firing and the fact that you are in pain. More on the explanatory gap in chapter two.

seeming, according to Levine and identity theorists. Since being in pain just is having your C-fibers fire, it is not possible the second is instantiated without the first. Nevertheless, there is no denying that it seems possible. This makes the determination of being in pain by having your C-fibers fire even more brute than the determination of badness by being an instance of cat burning.

So if the instantiation of being in pain is determined by the instantiation of having your C-fibers fire, we should conclude that it is brutally determined.⁵ If

(4) There is no brute determination

were true, though, being in pain couldn't be brutally determined. Consider this argument, just like a fragment of the argument against realism above:

(3*) If being in pain is actually instantiated, it is brutally determined.

(4) There is no brute determination. Thus,

(5*) Being in pain is not actually instantiated.

This is absurd—one of the premises must be false. Most people agree that being in pain is determined by neural or functional or behavioral properties. If so, the case for thinking it is brutally determined by neural or functional or behavioral properties is strong. (3*), then, is very plausible. So (4) is false. But if so, not only is the above argument unsound, so is (1)-(5) since it too includes (4).

I am not denying that brute determination is queer. If moral properties are brutally determined, then *ipso facto* they are queer. But we have reason to believe that some

⁵ Is this because we are assuming that it is instantiations of *neural* properties that determine instantiations of sensational properties? If we assume that instantiations of functional properties determine instantiations of sensational properties does the determination no longer seem brute? I doubt it. For any functional property, it seems perfectly sensible to ask why instantiations of it suffice for feeling any way at all. For any functional property, it seems possible that someone instantiates it but does not feel any way and for any functional property, it is conceivable that someone instantiates it but does not feel any way.

actually instantiated properties are queer in just that way and this reason has nothing to do with a belief in moral realism.

3. DISANALOGIES BETWEEN SENSATIONAL AND MORAL PROPERTIES

We started with an argument against realism, (1)-(5). It is unsound because (4) is false.

If it were true, no one would ever be in pain.

This is an extremely simple response to anti-realism. Perhaps it is oversimple. There are, of course, many disanalogies between being in pain and badness. In general, there are many disanalogies between sensational and moral properties. Maybe these make a difference.

Blackburn, for example, notes that it is *a priori* that if moral properties are instantiated, their instantiations are determined by the instantiation of natural properties. By contrast, it is only knowable *a posteriori* that if sensational properties are instantiated, their instantiations are determined by, say, the instantiation of neural properties. Blackburn says that this shows that the determination of moral properties is a sort distinct from the determination of sensational properties: it is determination_{a priori} (Blackburn (1993): 145-146).⁶

If so, then (1)-(5) is slightly off-track. Instead of arguing that if moral properties are instantiated, they are brutally determined and then denying that there is brute determination, the anti-realist might argue only that if moral properties are instantiated, they are brutally determined_{a priori} and then deny that there is brute determination_{a priori}.

⁶ Blackburn distinguishes sorts of supervenience not sorts of determination. However, nothing is left out of his reply by putting it as I have.

Sensational properties are not counterexamples to that last, weaker premise. They are not brutally determined_{a priori}.

Blackburn is right that we know *a priori* that if moral properties are instantiated at all, their instantiations are determined by instantiations of natural properties and that this marks a difference between the moral and sensational cases. We do not know *a priori* that instantiations of sensational properties are determined by, say, instantiations of neural properties. But this does not show that there are two sorts of determination. It shows only that we know about determination in two ways. Compare: that Blackburn is knowable *a priori* and that Blackburn is the actual author of *Essays in Quasi-Realism* is only knowable *a posteriori*. This does not show that there are two sorts of identity, identity_{a priori} and identity_{a posteriori}. It shows only that some identities are knowable *a priori*, some only *a posteriori*.

All my objection to

(4) There is no brute determination

needs is that sensational properties are brutally determined. My objection makes no claim about how we know this.

Another disanalogy between moral and sensational properties is noted by David Chalmers. Chalmers writes,

Moral facts are not phenomena that force themselves on us. When it comes to the crunch, we can deny that moral facts exist at all...The same strategy cannot be taken for [sensational] properties, whose existence is forced upon us. (Chalmers (1996): 83; I have replaced 'phenomenal' with 'sensational')

The idea is that no matter how implausible their instantiation might be, we have no choice but to admit that sensational properties are instantiated. Not so in the moral case.

If the instantiation of moral properties is implausible enough, we can deny that they are instantiated. Like Blackburn, then, Chalmers might claim that (1)-(5) is slightly off-track. The crucial claim should not be that there is no brute determination but, rather, that accepting that a property is brutally determined is a last resort. In the case of sensational properties, we are forced to the last resort. Being in pain is forced upon us. Not so badness. Given a choice between accepting that it is brutally determined and denying that it is instantiated, we should deny that it is instantiated.

If that is what he is saying, then Chalmers is right that it is possible to hold that moral properties are not instantiated. There is no disanalogy here with sensational properties, though. It is possible to hold that sensational properties are not instantiated. Someone who holds either thesis, it seems to me, holds a false thesis.

What needs arguing on Chalmers's behalf is that anti-realism is more plausible than realism (where realism includes that moral properties are brutally determined). Otherwise, accepting that moral properties are brutally determined is the last resort, not anti-realism. Of course, Chalmers does not argue that anti-realism is more plausible than realism. The quote comes from a book about mind, not metaethics. Lacking such argument, though, we have no reason to think the moral and sensational cases are importantly disanalogous.

We have reached a preliminary conclusion. (1)-(5) is unsound. Being in pain, like goodness and badness, is brutally determined. This suffices to show that

(4) There is no brute determination
is false.

4. GREEDINESS

There is a deeper response to the argument, namely, that not all moral properties are brutally determined and, hence,

(2) For any natural property, it is an open question *why* things instantiating that property instantiate a moral property, if they do

is false.

We earlier assumed that if badness is brutally determined, all moral properties are. Doing so helped to make (1)-(5) as strong as possible. The assumption, however, seems mistaken. To derive the conclusion that badness is brutally determined, we appealed to the thesis that for any natural property, it is an open question why things instantiating that property instantiate a moral property, if they do. The thesis does at least *seem* to be true of badness (although, as noted, some realists reject it).

Is a version of it true of all moral properties? Greediness is a moral property. Is a version of the thesis true of greediness? Say that Alan acts greedily. He and Bob work together, get paid, and Alan takes all the money for himself. The instantiation of being an instance of greediness here is determined by the instantiation of a natural property, being a taking it all for yourself. Is it really an open question why actions instantiating the second instantiate the first?

Consider someone who says, "I can see that what Alan did was a taking it all for himself but I just don't see why doing so was acting greedily." Isn't there something wrong with this person? He might lack the concept of greediness, in which case his question is like, "Why is something with four equal sides and four right angles a square?." Or he might lack information about the case. (Maybe he thinks that Bob is a dog. People who don't share money with dogs aren't greedy.) In that case, his question

is like “Why is a glass of H₂O a glass of water?.” Either way there is a striking difference between his question and, say, “Why is causing pain for fun bad?”

Thus, being greedy is a moral property but not brutally determined. I guess that all the virtue and vice properties are like that. So (2) is false. At best, only the instantiation of badness (and maybe goodness) is brutally determined.

It should be noted that the idea that moral properties are brutally determined is a bit surprising. Moral properties are evaluative properties. Consider some non-moral evaluative properties: being a good pen or being a good typist or...Is it at all plausible that instantiations of these properties are brutally determined? Take the question, “I can see that this pen is lightweight and grippable and leak-free and...but why is it a good pen?” Does someone who asks that question really have the concept of a pen (or really know what a leak is or...)? If you accept that the instantiations of some moral properties are not brutally determined and that the instantiations of some non-moral evaluative properties are not brutally determined, you have to wonder what is so special about goodness and badness such that they are brutally determined.

5. WHY ARE ANTI-REALISTS UNINTERESTED IN GREEDINESS?

We started with an argument against realism, (1)-(5), that proceeded from the premises that moral properties would be brutally determined and that there is no brute determination. We rejected the argument; it is unsound for two reasons. First, some actually instantiated properties are brutally determined and, second, some moral properties are not brutally determined.

Like the first, the second response is extremely simple but it is never even considered by the people who insist that the instantiations of moral properties are brutally determined. I can think of three reasons for this.

First, it might be said that being greedy is not a moral property and, hence, not an example of a non-brutely determined moral property. It might be said, in general, that virtue and vice properties are not moral properties. Here, it is worth noting that anti-realists tend to be consequentialists and that Moore himself is a consequentialist and a consequentialist might insist that the only moral properties are goodness, badness, rightness, and wrongness. I have no argument against this position but it is worth pointing out that it is a cost of endorsing (1)-(5) that you have to deny that virtue and vice properties are moral properties.

Second, it might be said that being greedy *is* brutally determined and, hence, is not an example of a non-brutely determined moral property. It might be conceded that there is a sense of 'being greedy' in which instantiations of it are not brutally determined. For example, that an act is 'greedy' can be explained by appeal to the fact that it is a taking it all for yourself but this sense of 'greedy' is non-moral. Only *bad* instances of taking it all for yourself are greedy in the moral sense and since instantiations of badness are brutally determined, so are instantiations of greediness.

This makes it seem as if badness is a part of an analysis of greediness, that greediness is badness *plus* something, for example, taking it all for yourself. There is no need to think of greediness that way. Philippa Foot's position, as I understand it, is that acting greedily is a way of acting badly but not that badness is a part of greediness. To

Foot, the relation between greediness and badness is like the relation between being scarlet and being red.⁷ Redness is no part of the analysis of scarletness.

Third, it might be said that greediness is a moral property, instantiation of which is not brutally determined, but, it might be further said, that if realism is true, goodness and badness are instantiated. Since instantiations of goodness and badness would be brutally determined and imperceptible, the fact that instantiations of greediness would not be is irrelevant to the falsity of realism. It just shows that the argument needs to be modified slightly. The thrust remains. If realism were true, goodness and badness would be instantiated. They aren't, so it isn't.

The problem with this maneuver is seeing why realism has to include the thesis that goodness and badness are instantiated. The minimal realist thesis is just that mind-independent moral properties are actually instantiated. That goodness and badness are among them is an optional extra. Some realists deny that goodness and badness are properties at all, much less moral properties. Peter Geach and Judith Jarvis Thomson are examples.⁸

6. CONCLUSION

What has emerged is that even against a sort of realism that accepts that moral properties are brutally determined, (1)-(5) is unsound and unsound for a simple reason. If sound, it would show that no one is ever in pain. Since my head hurts right now, the argument is unsound. We have seen no reason to give up this simple reason. Furthermore, it is worth

⁷ Cf. Foot (2000). Her position seems also to be that instantiations of goodness are determined by instantiation of virtue properties and the determination is not brute. Since virtue properties are not brutally determined, she thinks that goodness is not brutally determined.

⁸ Cf. Geach (1956) and Thomson (1997).

stress that it is only a *sort* of realism that accepts that all moral properties are brutally determined; plenty of realists deny that they all are.

Regardless, there is something attractive about (5) and, thus, something attractive about the thesis that realism is false. These theses have wide, intuitive appeal and showing that two arguments for them are unsound is not going to lessen it.

The appeal must come from thinking that moral properties are not the sort of thing science investigates and then supposing it to be hard to see how they relate to those properties science does investigate. “There is no room,” Mackie writes, “for objective moral [properties]...in this strictly materialist universe” (Mackie (1980): 7; I have replaced ‘qualities’ with ‘properties’).

Maybe it is hard to see where the room is. Still that isn’t sufficient reason to think there is no room. Also, ‘strictly materialist’ is misleading. Even before ‘strict materialism’ was the scientific worldview, people thought there was no room for moral properties. Hamlet was not endorsing strict materialism when he said that nothing is good or bad but thinking makes it so. If there is good reason to doubt realism, it was visible to a Danish prince who believed in souls, ghosts, and the science of the late middle ages.

CHAPTER TWO THE EXPLANATORY GAP

Take a world that is physically just like this one but in which no one is conscious. This is a zombie world. In the actual world, I am pricked with a pin, I have a sensation of pain, and flinch. In the zombie world, I am pricked with a pin and flinch. There is no sensation.

It matters whether there could be a zombie world. If there could, there could be a world just like this one physically but different sensationally. So sensational properties—being in pain, being in ecstasy, being in agony, etc.—are not physical properties and sensational states are not physical states. Analogous considerations would show that sensational properties are not functional properties and sensational states are not functional states.

There is a familiar, simple argument that there could be a zombie world: it is conceivable that there is one and, hence, it is possible that there is one. Most people think there is something wrong with the argument—it can't be that easy to show that sensational properties are not physical properties.

Joseph Levine offers two objections to the argument. I reject the first but what is wrong with it is interesting: it yields some constraints on a successful objection. I then reject the second but what is wrong with it is also interesting: it shows that a certain theory of mind is analogous to a certain metaethic.

First, more on the familiar, simple argument.

1. THE ZOMBIE ARGUMENT

We can put the argument—I call it “the zombie argument”—like this:

- (C1) It is conceivable that there is a zombie world.
- (C2) If a proposition is conceivable, it is possible. Thus,
- (C3) It is possible that there is a zombie world.

A proposition is conceivable just in case (i) it is possible to imagine a situation that verifies the proposition and (ii) the imagining is clear and distinct (in a sense to be presently explained).

Because of (i), it is insufficient for conceiving that this table is a bit to the right that you cannot rule out that it could be. To conceive that proposition, you have to go through the process of imagining a situation in which the table is a bit to the right. By contrast, when I say, “It is conceivable that Nader wins the election,” I have not imagined a situation in which he does; I never have imagined such a situation. I am just voicing my belief that I have not ruled out that it is possible or, more strongly, voicing my belief that it is possible.

Because of (ii), it is insufficient for conceiving that this table is a bit to the right that it seems, off the top of your head, that you can imagine a situation in which it is. Take the proposition that Alan constructs a regular, nineteen-sided polygon with just a ruler and compass. Off the top of my head, it seems I can imagine a situation verifying this. On rational reflection, I cannot. Each attempt is a failure. (I take David Lewis’s word for this (Lewis (1986): 90).) Each attempt turns out to be imagining a situation in which Alan constructs a regular, *seventeen*-sided polygon or in which Alan is merely *trying* (failing) to construct a regular, nineteen-sided polygon or... The “clear and distinct” clause in (ii) is meant to limit what you can conceive to those propositions

which, on rational reflection, you can imagine being verified. This makes (C2) more plausible than it would be without the clause.

How rational? How much reflection? These are good questions but answering them is not important to what follows. Just rational reflection or rational reflection after learning a lot of empirical facts? More on that later.

If it is possible to, on rational reflection, imagine a situation verifying that there is a zombie world, then

(C1) It is conceivable that there is a zombie world

is true. Not everyone will accept that this is possible and much more needs to be said about clear and distinct imagining but it is *prima facie* plausible that I can imagine such a situation. Since I am interested in objections to (C2), from now on I just assume (C1).

“Possible” in

(C2) If a proposition is conceivable, it is possible

means *metaphysically possible*. So (C2) is quite strong for two reasons: first, because it claims that conceivability *entails* possibility and, second, because the possibility in question is metaphysical. More plausible than conceivability *entailing* possibility is the weaker thesis that if a proposition is conceivable, that is good, defeasible evidence that it is possible. We do take conceivability to be good, if defeasible, evidence for possibility. Reading some works of fiction, you conceive their content and come to believe, defeasibly, that things could be so. Conceiving that I am taller, I come to believe, defeasibly, that I could be taller. It is less plausible that whatever is conceivable is possible. For example, I seem to be able to conceive that I exist without my parents ever existing and that proposition is impossible. Friends of (C2) deny that I can conceive this

impossibility but because it at least seems as if I can, it is more plausible that conceivability is merely good, defeasible evidence for possibility than that it entails possibility.

Also, if conceivability entails any sort of possibility, it is more plausible that it entails epistemic possibility than that it entails metaphysical possibility. A proposition is epistemically possible just in case it is not *a priori* false (I take this (slightly non-standard) terminology from Levine (1993)). It is somewhat plausible that if I can imagine, on rational reflection, a situation verifying a proposition, then that proposition is not *a priori* false. To see exactly how plausible, more must be said about rational reflection. At any rate, the apparent counterexample to the thesis that conceivability entails metaphysical possibility does not apply to the thesis that conceivability entails epistemic possibility. That I exist without my parents ever existing is not *a priori* false, hence, it is epistemically possible. So the fact that that proposition is conceivable—if it is a fact—would not show that conceivability does not entail epistemic possibility. It would only show that it does not entail metaphysical possibility.

I grant that (C2) is less plausible in both of the above ways than some alternatives. Still, it is the premise Levine objects to and it is the premise most often used along with (C1) to derive

(C3) It is possible that there is a zombie world.

From (C3), it follows that sensational properties are not physical properties and sensational states are not physical states. What, if anything, is wrong with this argument?

2. LEVINE'S FIRST OBJECTION

Levine suggests, roughly, that

(C2) If a proposition is conceivable, it is possible

is false because what is conceivable is an epistemic matter, what is possible is a metaphysical matter, and there is no entailment from epistemic premises to a metaphysical conclusion.⁹ This presentation of the objection is rough because metaphysical conclusions follow from certain epistemic premises. From the premise that you know you are a philosopher, it follows that you are a philosopher. This is a metaphysical conclusion, entailed by an epistemic premise. So the no-metaphysics-from-epistemology objection has to be precisified, the rough statement of it is false.

Levine offers the following: a proposition is conceivable just in case it is epistemically possible. A proposition is epistemically possible just in case it is not *a priori* false. Thus, a proposition is conceivable just in case it is not *a priori* false. For example, that water is XYZ, not H₂O is conceivable since it is not *a priori* false. However, that proposition is metaphysically impossible. So some conceivable propositions are impossible. So (C2) is false and (C1)-(C3) is unsound.¹⁰

This is recognizably a precisification of the no-metaphysics-from-epistemology objection. Levine's claim is that no conclusions about possibility follow from certain epistemic premises, those asserting that a proposition is not *a priori* false.

⁹ "[The problem with (C1)-(C3) is its] using epistemological premises to support a metaphysical conclusion" (Levine (1993): 124). On the previous page, he writes, "What is imaginable is an epistemological matter and therefore, what imagining that [there is C-fiber firing without pain] does is establish the epistemological possibility that pain is not identical to the firing of C-fibers. It takes another argument to get from the epistemological possibility that pain is not the firing of C-fibers to its metaphysical possibility..." (123)

¹⁰ See Levine (1993): 123. Levine gives a much more complex version of this response in Levine (2001). The response is more complex, however, only because the version of the zombie argument is. The main idea persists.

3. WHAT IS WRONG WITH LEVINE'S FIRST OBJECTION

Levine is right that there is a use of “conceivable” according to which a proposition is conceivable just in case it is not *a priori* false. That sort of conceivability is not at issue in the zombie argument, though. This emerges as follows.

I know that I could be a little taller. If Levine’s objection to (C2) is sound, that I could be a little taller is not *entailed* by the premise that it is conceivable that I am a little taller. That premise should at least, I think, be evidence for the conclusion. As I mentioned earlier, we standardly do take it to be evidence. If the objection to (C2) is sound, though, it is not. How could it even be evidence for the proposition that I could be a little taller that *a priori* I cannot tell if the proposition is false? Of course I can’t tell *a priori*. That just seems irrelevant to whether I could be a little taller.

Or consider this case: I am wondering what the atomic number of gold is. *A priori*, it could be one, two, three,... Not only does that not entail that the atomic number of gold could be one, two, or three, neither does it provide any evidence that one or two or three could be the atomic number of gold. (Maybe it provides a *tiny* bit of evidence? I doubt it. Say that your empirical evidence equally supports that the atomic number of gold is two and that it is three. You are looking for a way to decide between the two. If you find that *a priori* it is not false that it is two, do you increase *at all* the probability that it is two? And then, finding that it is also not *a priori* false that the atomic number is three, inch up its probability?)

The first problem with the objection to (C2), then, is that it is too strong. It would show not only that conceivability does not entail possibility but, also, that it is no

evidence for possibility in cases where it ought to be (and is assumed to be). Put aside for a moment whether

(C1) It is conceivable that there is a zombie world

is evidence for

(C3) It is possible that there is a zombie world.

Regardless of whether conceivability is evidence for possibility in this controversial case, it at least seems to be evidence in other cases. That it is conceivable that I am a little taller seems to be evidence that I could be a little taller. Likewise, that it is conceivable that there is a behavioral duplicate of me that is not a mental duplicate of me is evidence that there could be such a duplicate and, hence, that behaviorism is false.

There is a second problem with Levine's objection. Consider an argument in *Naming and Necessity* (Kripke (1980)) that is like (C1)-(C3). The argument is:

(K1) It seems possible that there is a zombie world. Thus,

(K2) It is possible that there is a zombie world (absent something that defeats the seeming possibility).

Kripke then points out that he cannot think of a defeater. (He does not say there is none.)

Nevermind what Kripke thinks, let us infer

(C3) It is possible that there is a zombie world.¹¹

(K1)-(C3) is different from, but related to, (C1)-(C3), of course. Obviously, their conclusions are the same. Also, one reason to think that (K1) is true is because (C1) is true. Conceivability provides evidence for seeming possibility.

¹¹ There are two ways in which (K1)-(C3) differs from the argument in Kripke (1980) that I should mention. First, (K1)-(C3) puts the argument in terms of zombie worlds; the argument in Kripke (1980) is put in terms of zombie individuals. Second, more stress in Kripke (1980) is put on the seeming possibility of sensational properties being instantiated without physical properties than on the seeming possibility of physical properties being instantiated without sensational properties.

Levine's response to (C2) cannot be modified to respond to (K1)-(C3). Levine claims that all propositions that are not *a priori* false are conceivable but it is not the case that all propositions that are not *a priori* false seem possible. For example, it does not seem possible that water is XYZ. This is a problem for Levine's objection: it cannot respond to an argument similar to (C1)-(C3) and with the same conclusion.

So there is something wrong with Levine's objection. Since its second premise—some epistemically possible propositions are metaphysically possible—is true, his first premise—all epistemically possible propositions are conceivable—is not. If it were true, conceivability wouldn't even be a guide to possibility (or seeming possibility) and, anyway, the objection only applies to (C1)-(C3), not its relative, (K1)-(C3).

Two constraints on an objection to the zombie argument have emerged. A successful objection has to be in one way less strong than the no-metaphysics-from-epistemology objection. It has to allow that conceivability is evidence for possibility in humdrum cases. For example, it has to allow if it is conceivable that I am a little taller, that is evidence that I could be a little taller.

A successful objection also has to be in one way less weak than the no-metaphysics-from-epistemology objection. It (or a modification of it) has to respond to (K1)-(C3).

4. THE EXPLANATORY GAP

Levine has another objection to (C1)-(C3), one which meets both constraints. A lot of background is needed to see it.

For simplicity, let us assume that if it is a physical property at all, then being in pain is the property of having your C-fibers fire. If being in pain is having your C-fibers fire, then if it is a fact that your C-fibers are firing, then it is a fact that you are in pain. However, the fact that your C-fibers are firing seems to be something else entirely from the fact that you are in pain. Levine thinks that there is an *explanatory gap* between the two facts. When he introduced the phrase “explanatory gap,” he illustrated it like this:

There seems to be nothing about C-fiber firing which makes it naturally ‘fit’ the phenomenal properties of pain, any more than it would fit some other set of phenomenal properties...[T]he identification of the qualitative side of pain with C-fiber firing...leaves the connection between [them] completely mysterious. (Levine (1983): 356)¹²

The gap between the fact that your C-fibers are firing and that you are in pain is an *explanatory gap* because if we had an explanation for the obtaining of one in terms of the obtaining of the other, the connection between the two facts would not be mysterious.

We have a rough definition:

EXPLANATORY GAP—There is an explanatory gap between two facts iff it is mysterious how to explain why the (would-be) explanandum obtains in terms of why the (would-be) explanans does. (From now on, I drop the ‘would-be’s.)¹³

¹² Cf. Block (1994): “This is the famous explanatory gap. We have no conception of our physical or functional nature that allows us to understand how it could explain our subjective experience...Francis Crick and Christoff Koch have famously hypothesized that the neural basis of consciousness is to be found in certain phase-locked 40 Hz neural oscillation. But how does a 40 Hz neural oscillation explain *what it’s like*...to be us? What is so special about [that] oscillation as opposed to some other physical state?” (210)

¹³ The use of “mysterious” suggests that the existence of an explanatory gap is relative to thinkers and times. Nothing is mysterious *simpliciter*. Some things are completely mysterious to me now but will not be later. Some things will always be mysterious to me but not to a physicist. Levine does not say whether he thinks explanatory gaps are relative to thinkers and times. I think he should.

The attitude of some philosophers towards the explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain is that there is for all people, at present, a gap between them, but, in the future, there will not be. The attitude of other philosophers is that for all people, now and for always, there is a gap but for other, smarter creatures there might not be. The attitude of still other philosophers is that there is a gap for all people at present and we just do not know if there always will be or whether there would be such a gap for anyone.

However, for simplicity, I drop the reference to thinkers and times when talking about explanatory gaps.

This is rough because 'mysterious' is itself pretty mysterious. To clear up the notion of an explanatory gap, it will help to look at some pairs of facts between which there are no gaps.

There is no explanatory gap between the fact that something has just four equal sides and four right angles and the fact that it is a square. That the first obtains explains why the second does. It follows from the meaning of "square" that anything with just four equal sides and four right angles is a square.

Likewise, there is no explanatory gap between the fact that the glass is full of H₂O and the fact that it is full of water. Unlike the square case, though, the explanation is not available *a priori*. *A priori* it is mysterious how the obtaining of the fact that the glass is full of H₂O explains the obtaining of the fact that it is full of water. The explanation comes, *a posteriori*, from chemistry.

So Levine's claim that there is an explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain is not just that it is mysterious *a priori* how the explanation would go. His claim is that it is mysterious, period. No matter how much neuroscience you know, how much psychology, etc., it is mysterious.

We can distinguish two dimensions of mysteriousness: degree and valence. There are two degrees of mysteriousness: total and partial. Levine is clear that he thinks the gap between the fact that there is a C-fiber firing and the fact that there is a pain is only partial. He thinks that what needs explaining when the fact that you are in pain obtains is complex: that you are feeling a certain way, that you are disposed to act a certain way, and that you are disposed to desire certain things, among others. The obtaining of the fact that your C-fibers are firing, Levine thinks, can explain the obtaining of the second

two facts. What it cannot explain is why the fact that you are feeling a certain way (“the qualitative side of pain”) obtains.

An attempt to explain why the fact that you are hurting obtains in terms of the obtaining of the fact that your C-fibers are firing, Thomas Nagel writes, would be like a pre-Socratic trying to explain the obtaining of the fact that something is matter in terms of the obtaining of the fact that it is energy (Nagel (1979): 177). In the pre-Socratic case, however, unlike in the pain/C-fiber case, the mysteriousness is total. Nothing about the obtaining of the fact that something is matter can be explained by the pre-Socratic in terms of the obtaining of the fact that it is energy.

Another dimension of mysteriousness is its valence. We can distinguish three sorts. When the explanandum and explanans seem to oppose one another, this makes it mysterious how the obtaining of one is to be explained in terms of the obtaining of the other. The facts (seem to) have opposed valences. For example, consider an early modern view of the relation between the fact that something is extended and the fact that it is colored. According to it, the obtaining of the first explains why the second *does not* obtain.¹⁴ Presented with the contemporary physicalist view that colors are surface properties of bodies, the early modern would be as mystified as the pre-Socratic—the mystery would be total. The reason he would be mystified, though, is that he has a theory according to which the fact that something is extended entails that it is not colored. He thinks the fact that something is extended opposes the fact that it is colored.

In the pre-Socratic case, the fact that something is energy and the fact that something is matter do not seem to oppose each other, they seem indifferent to each

¹⁴ Richard Price writes, “A *coloured body*, if we speak accurately, is the same absurdity with a *square sound*. We need no experiments to prove that...colours...are not real qualities of bodies; because the ideas of matter and of [colours] are incompatible.” (Price (1969): 680-681)

other. The facts have neutral valences. When the explanans and explanandum seem indifferent to each other, this makes it mysterious how the obtaining of one is to be explained in terms of the other.

Finally, it is possible to have an explanatory gap between two facts even when their valences are attractive. This can happen when the explanans seems to support the explanandum but only partially. In a case like this, it is mysterious how the obtaining of one is to be explained in terms of the obtaining of the other—something seems left out. This is Levine's position: the fact that your C-fibers are firing partially supports the fact that you are in pain—that is why the mystery is not total—but the support is only partial since the feel of pain is unexplained.

Summing up, there is an explanatory gap between two facts just in case it is mysterious how to explain the obtaining of the explanandum in terms of the obtaining of the explanans. We have distinguished two dimensions of mysteriousness. Some mysteries are total, some partial. Some mysteries arise because the two facts seem opposed to each other, some because the facts seem indifferent to each other, and others because the facts seem supportive but only partially. Levine's position is that there is an explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain. The mystery is partial and it arises because the facts seem only partially supportive.

5. LEVINE'S SECOND OBJECTION

The existence of explanatory gaps is not surprising and the world is full of them. There is an explanatory gap between the fact that someone is Joseph Levine and the fact that my car is a Chevy. It *is* mysterious how to explain why the second obtains in terms of why

the first does. All that shows is that the fact that someone is Levine is a bad choice of explanans for the fact that my car is a Chevy. To explain why the fact that my car is a Chevy obtains, you are better off appealing to the fact that it came from a certain factory and was built out of certain materials and according to certain specifications.

Levine claims that if there is an explanatory gap between two facts, it is conceivable that one obtains without the other, in particular, it is conceivable that the explanans obtains without the explanandum. In "Materialism and Qualia," for example, he writes,

If there is nothing we can determine about C-fiber firing that explains why it has the qualitative character it does...it immediately becomes imaginable that there be C-fiber firings without the feeling of pain, and vice versa. (Levine (1983): 357-358; Levine uses 'conceive' and 'imagine' interchangeably)

And in "On Leaving Out What It's Like," he writes,

It is because the qualitative character is itself left *unexplained* by the physicalist...theory that it remains conceivable that a creature should occupy the relevant physical...state and yet not experience qualitative character. (Levine (1993): 129)

Since there is a gap between the fact that someone is Levine and the fact that my car is a Chevy, it is conceivable that someone is Levine and it is not the case that my car is a Chevy. There is a gap, Levine thinks, between the fact that your C-fibers are firing and the fact that you are in pain, so it is conceivable that your C-fibers are firing but you are not in pain. Levine thinks, further, that there is an explanatory gap between any physical fact and any sensational fact. Thus, he thinks that

(C1) It is conceivable that there is a zombie world

is true.

Levine also thinks that there can be an explanatory gap between two facts even though that one obtains implies that the other does. What is explicable, he thinks, is one thing, what implies what is something else entirely.

Take the gap between the fact that your C-fibers are firing and the fact that you are in pain. Say that C-fiber firing is pain and, thus, that your C-fibers are firing implies that you are in pain. This does not close the explanatory gap. However, it makes the gap—unlike the gap between the fact that someone is Levine and the fact that my car is a Chevy—surprising. Of course there is an explanatory gap between the fact that someone is Levine and the fact that my car is a Chevy—those facts have nothing to do with each other. Levine thinks there can be gaps between facts that have an awful lot to do with each other. There can be *objectionable explanatory gaps*.

OBJECTIONABLE EXPLANATORY GAP—There is an objectionable explanatory gap between two facts iff there is an explanatory gap between them and that the explanans obtains implies that the explanandum does.^{15,16}

¹⁵ There is an objectionable explanatory gap between the fact that water is H₂O and the fact that twice two is four. There is no explaining why the second obtains in terms of why the first obtains but since both facts obtain in all possible worlds, each implies the other. From now on, I ignore objectionable explanatory gaps between facts that obtain in all possible worlds.

¹⁶ Levine's most famous thesis is that there is an objectionable explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain. As a result, he thinks that even if you know that when one obtains, the other does, it makes sense to wonder why this is so.

More controversially, Levine's view is that even if pain is C-fiber firing, there can be an explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain. Some deny this. They say that if pain is C-fiber firing, it is about as sensible to wonder why the obtaining of the fact that someone's C-fibers are firing explains the obtaining of the fact that he is in pain as it is to wonder why the obtaining of the fact that someone is Levine explains the obtaining of the fact that he is Levine. (Papineau (1996), for example, puts the objection clearly.)

Admittedly, there is no explanatory gap between the fact that someone is Levine and the fact that someone is Levine. That is insufficient to show that there is no explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain. That Levine is Levine is knowable *a priori* does not show that the identity of pain and C-fiber firing is or the identity of water and H₂O.

The case of water and H₂O is instructive. As long as you don't know any chemistry, it can make good sense to wonder why the fact that something is a glass full of H₂O explains why it is a glass full of water. Further, we can give a perfectly good explanation.

We have a spectrum of cases here. There is no explanatory gap between the fact that someone is Levine and the fact that someone is Levine. There is just no explanation to be offered here. There is also no explanatory gap between the fact that there is a glass full of H₂O and the fact that there is a glass full of

We now have an argument against

(C2) If a proposition is conceivable, it is possible.

We saw above that Levine commits himself to

(L1) If there is an explanatory gap between two facts, it is conceivable that the explanans obtains without the explanandum.

Take two facts between which there is a gap.

(L2) There is an explanatory gap between x (explanans) and y (explanandum).

From (L1) and (L2), a preliminary conclusion follows,

(L3) That x obtains and y does not is conceivable.

Levine thinks that objectionable explanatory gaps are possible. So we are free to stipulate that the gap between x and y is objectionable.

(L4) The gap between x and y is objectionable.

So that x obtains and y does not is impossible. From (L3) and (L4), it follows that

(L5) That x obtains and y does not is conceivable and impossible.

From which follows a final conclusion,

(L6) (C2) is false.

Levine never explicitly offers (L1)-(L6). His explicit objections to (C1)-(C3) are all versions of the no-metaphysics-from-epistemology objection. He holds the crucial premises of (L1)-(L6), though, so I assume he would endorse the argument.

There is a lot to like about it. Unlike the no-metaphysics-from-epistemology objection, it is not too strong. It allows that the conceivability of a proposition is

water but this is not because there is no explanation to be offered. Rather, it is because we have a good explanation. That is just what is missing in the mental case and why there is an explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain.

evidence for its possibility. For all it says, conceivability entails possibility *except in cases involving objectionable explanatory gaps*. If that is right, conceivability would be a good, defeasible guide to possibility.

Unlike the no-metaphysics-from-epistemology objection, this second objection to (C2) is not too weak. It can be adapted to respond to Kripke's argument, (K1)-(C3). Not only is

(L1) If there is an explanatory gap between two facts, it is conceivable that the explanans obtains without the explanandum

true, so is

(L1*) If there is an explanatory gap between two facts, it seems possible that the explanans obtains without the explanandum.

Like Levine, someone who endorses (K1)-(C3) accepts that there is an explanatory gap between the fact that your C-fibers are firing and the fact that you are in pain. So, from (L1*), it follows that it seems possible that your C-fibers are firing but you are not in pain. Someone who endorses (K1)-(C3), however, insists that this is not a mere seeming possibility. It is a genuine possibility.

But, Levine can point out, because of the explanatory gap, it would seem possible whether or not it was possible and, further, there is a defeater of the seeming possibility: the physicalist case for thinking that being in pain is a physical property. The physicalist case alone would not do this. Someone who endorses (K1)-(C3) might do so in spite of accepting that there is a strong physicalist case. This is the textbook reading of Kripke, I think. On the textbook reading, Kripke accepts

(C3) It is possible that there is a zombie world

but acknowledges that there is a strong case against it. Levine is in position both to explain why that your C-fibers are firing but you are not in pain *seems* possible—because of the explanatory gap—but is impossible.

There were two problems with the no-metaphysics-from-epistemology objection. These led to two constraints on a successful objection to (C2). Levine's second objection, (L1)-(L6), meets both.¹⁷

6. WHAT IS WRONG WITH LEVINE'S SECOND OBJECTION

(L1) is not true, though. There are counterexamples to it. One comes out of the following case.

You are looking at a painting. You do not know whether it is beautiful. You know that you like it, but I distinguish that from knowing whether it is beautiful. I hope this is not controversial. Plenty of people like paintings they know are ugly.

A friend helpfully points out that the painting has certain features: it is composed a certain way and uses certain colors. You then see that it is a beautiful painting.

¹⁷ Levine endorses

(L1) If there is an explanatory gap between two facts, it is conceivable that the explanans obtains without the explanandum

and also something like the converse:

If it is conceivable that an explanans obtains without an explanandum, there is an explanatory gap between the two.

"Something like" because, stated that way, the converse is not true. Why did the egg break? It fell. That is a fine explanation, there is no explanatory gap here, but it is conceivable that the egg falls and not break. So the converse thesis needs to be precisified.

Take the imprecise thesis, though. Combined with (L1), we get a biconditional:

There is an explanatory gap between two facts iff it is conceivable that the explanans obtains without the explanandum.

I doubt Levine endorses the biconditional as an explanation of what an explanatory gap is. (See, for example, his replies to critics at the 2004 Pacific APA.) Still, if you take the biconditional to be explanatory, you have a clearer explanation of the explanatory gap than the one I offered. Unlike that attempt, the biconditional does not make use of the mysterious 'mysterious.' This is further reason to like (L1): it helps provide a less mysterious account of the explanatory gap than the one Levine and I give.

Furthermore, you come to believe that any painting that looked that way would be beautiful.

Still, it might remain mysterious to you why paintings that look like that are beautiful. Satisfied *that* the painting is beautiful, you might nevertheless ask your friend *why* the obtaining of the fact that the painting is beautiful is explained by the obtaining of the fact that it looks a certain way. This suggests that there is an explanatory gap between the two facts. (Levine is satisfied *that* pain is C-fiber firing, he just wants to know *why* the obtaining of the fact that your C-fibers are firing explains the obtaining of the fact that you are in pain.)

If (L1) is true, then, it is conceivable that the painting looks just the way it does but is not beautiful. That is not conceivable, though. What is conceivable is that the painting looks just that way and you do not like it and, also, that it looks just that way and other people do not like it. It is also conceivable that it looks just that way and you and others do not *think* it is beautiful. That it looks that way and is not beautiful, though, is inconceivable. So

(L1) If there is an explanatory gap between two facts, it is conceivable that the explanans obtains without the explanandum

is false.

There is another counterexample to it. It is provided by a Moorean metaethic. Consider three Moorean theses—whether they are exactly theses of *Principia Ethica* (Moore (1993)), I do not argue. They are at least like theses in that book.

First, badness *simpliciter* is a moral property. Burning cats just for fun, for example, is bad. Not just bad for your conscience or bad for the cats. Bad, period.

Second, if any moral property is instantiated, its instantiations are determined by the instantiation of non-moral properties. The instantiation of a property *x* is determined by the instantiation of a property *y* just in case the instantiation of *y* suffices for the instantiation of *x* and this is true across all possible worlds, not just the actual world. Since badness is a moral property, if it is instantiated, its instantiation is determined by the instantiation of non-moral properties. Being an instance of burning a cat just for fun, for example, is a non-moral property. Say that Alan burns a cat just for fun (for short, Alan burns a cat). His act is bad and the instantiation of badness here is determined by the instantiation of the non-moral property, being an instance of cat burning.

Third, for any non-moral property, it is an open question *why* things instantiating that property instantiate badness, if they do. It is an open question *why*, for example, when someone burns a cat, his act is bad. It is not an open question *whether* his act is bad. Everyone knows that it is. What is open is *why* this bad act is a bad act.

Moore does not say much about what an open question is. An example will fix ideas about it. Take the question, “Why is cat burning bad?” Take, also, a conceivable answer to it: cat burning is bad because it hurts the cat. Someone might accept the answer—I do—without thinking it is a perfectly satisfying explanation. The answer just produces a new question, “Why is hurting cats bad?” Take a conceivable answer to that question: hurting cats, and, thus, burning cats, is bad because cats don’t like it. Someone might accept *that* answer while still thinking it, too, is a less than perfectly satisfying explanation. He might wonder, “Why is doing something that cats don’t like bad?” A trail of why-questions has been opened up and it is hard to see where (or if) it stops.

Each answer to each question in the series seems to raise a question analogous to the first. This is the sense in which it is an open question why cat burning is bad.¹⁸

Moore, of course, did not talk about explanatory gaps. However, just as for any non-moral property, it is an open question why something instantiating it instantiates badness (if it does), I think Moore would have accepted that there is an explanatory gap between any non-moral fact and the fact that an action is bad.

Hence, since it is an open question why cat burning is bad, it is safe to assume that there is an explanatory gap between the fact that an action is a cat burning and the fact that it is bad. If there were no such gap, there would be an explanation for why cat burning is bad and if there were such an explanation, it would not be an open question why it is bad.¹⁹

If

(L1) If there is an explanatory gap between two facts, it is conceivable that the explanans obtains without the explanandum

is true, it is conceivable that an action is a cat burning but it is not bad. This is inconceivable, though. Perhaps it is conceivable that people think cat burning is good or are all for cat-burning but it is inconceivable that cat burning is not bad.

That it is inconceivable is crucial to Moorean moral epistemology. The way you figure out which things are bad is by forming a clear and distinct idea of them and seeing if they strike you as bad. There is no clearly and distinctly conceiving that there is a cat burning without having it strike you as bad.

¹⁸ Much more on open questions is in chapter one.

¹⁹ Note that on Moore's view, the gap is objectionable. That an action is a cat burning implies that the action is bad because of Moore's determination thesis: the instantiation of being an instance of cat burning determines the instantiation of badness. In this way, the gap between the fact that an action is a cat burning and the fact that it is bad is like the gap between the fact that your C-fibers are firing and the fact that you are in pain.

We thus have another counterexample to (L1). Well, calling it a “counterexample” might be pushing it. It is a counterexample if the three Moorean theses above are possible. There are reasons to think they are not.²⁰ Even if they are not, though, is it not plausible that this is because they conflict with (L1).

7. DIFFERENCES BETWEEN BEING IN PAIN, BEAUTY, AND BADNESS

There are a lot of differences between the pain and beauty and badness cases, of course. Maybe they matter to whether the beauty and badness cases are genuine counterexamples to (L1). I focus on disanalogies between beauty and pain. There are similar disanalogies between badness and pain.

There is an explanatory gap between the fact that a painting looks a certain way and the fact that it is beautiful. However, it is *a priori* that if any painting is beautiful, then it looks some way or other. By contrast, take the gap between the fact that your C-fibers are firing and the fact that you are in pain. It is not *a priori* that if you are in pain, then you are in some physical state or other. It might be true but it is not *a priori*.

This is a genuine difference between the two cases but it does not have anything to do with whether there is an explanatory gap in both cases. At best, I think, it shows that the degree of mysteriousness between the fact that the painting is beautiful and the fact that it looks a certain way is partial. You know *a priori* that beautiful paintings have to look some way or other so it is not totally mysterious why the obtaining of the fact that the painting is beautiful is to be explained by the obtaining of *some fact or other* about how it looks. What you don't understand is why the obtaining of the fact that the painting looks *this* way explains the obtaining of the fact that it is beautiful. This is like

²⁰ See Thomson (1997) which argues that there is no property badness, much less a moral property.

Levine's position. He knows—albeit *a posteriori*—that the obtaining of the fact that you are in pain is to be explained by the obtaining of *some physical fact or other*. What he does not understand is why the obtaining of the fact that your C-fibers are firing explains the obtaining of the fact that you are in pain.

A further difference between the pain and beauty cases is that whereas it is clear in the pain case what needs explaining—the way it feels—it is not clear in the aesthetic case. That beautiful things merit appreciation? That they produce pleasure? That they produce pleasure in those who *know* they are beautiful?

Whereas we have a pretty good grip on what pain is, we have a pretty bad grip on beauty with the result that it is not so clear what needs explaining in the beauty case. Again, this is a genuine difference but, again, I do not think it shows that there is a significant difference between the beauty gap and the pain gap. That we do not know what is missing from our explanation of why the fact that that painting is beautiful obtains does not show that nothing is missing.

A third difference between the pain and beauty cases is that it does not seem possible that just this appearance fact obtains and the fact that the painting is beautiful fails to obtain; by contrast, it does seem possible that the fact that your C-fibers are firing obtains and the fact that you are in pain does not. Maybe this shows that the idea of an explanatory gap in the beauty case is out of place. You only get an explanatory gap when it seems like the explanans could obtain without the explanandum. If so, the beauty case is no counterexample to (L1). There is no explanatory gap between the fact that a painting looks this way and the fact that it is beautiful.

Again, there is a genuine difference here but, again, it does not show that there is no explanatory gap in the beauty case. The idea of an explanatory gap between two facts is that there is no explaining why the one obtains in terms of why the other does. We seem to have that in the beauty case. That it does not seem possible that the appearance fact obtains without the beauty fact does not show that there is no gap unless we insist on the thesis that if there is an explanatory gap between two facts, then that the explanans obtains without the explanandum seems possible. Like (L1), I think this thesis is plausible but false. The beauty case is a counterexample.

Finally, most roughly, it might be objected that there is a total mystery about why the obtaining of the fact that your C-fibers are firing explains the obtaining of the fact that you are feeling any way at all. There is an aspect of being in pain that is left totally unexplained by the fact that your C-fibers are firing. In the beauty case, it might be objected, things are otherwise. Nothing is totally mysterious. You can see that the obtaining of the appearance fact has *something* to do with the obtaining of the beauty fact. This is shown by the fact that you saw that the colors and composition of the painting were relevant to its being beautiful. Only a partial mystery remains.

I do not know whether the objection is right but, even conceding that it is, this difference does not show that there is no explanatory gap in the beauty case. It just shows that the gap in the beauty case differs slightly in character from the gap in the pain case. Both gaps are partial but the mental gap is partial because one aspect of being in pain is left totally mysterious.²¹ According to

²¹ This is a difficult issue. As mentioned above, it might be that part of being beautiful is meriting appreciation. If so, then perhaps there is a total mystery about how to explain the fact that the painting merits appreciation in terms of the fact that it looks a certain way.

(L1) If there is an explanatory gap between two facts, it is conceivable that the explanans obtains without the explanandum,

there is a connection between explanatory gaps and conceivability, not just a connection between explanatory gaps of a certain character and conceivability.

So I conclude that despite the disanalogies between pain and beauty, the beauty case is a counterexample to (L1). *Mutatis mutandis* for badness. This will be controversial. It depends, among other things, on separating a painting's being beautiful from its being liked and that is controversial and I haven't argued for it. Likewise, the badness case is only a counterexample if the three Moorean theses are possible and that is controversial (see footnote eleven) and I have not argued for it.

Whether or not the beauty and badness cases are counterexamples, though, they have a structure that shows that (L1) is false. They are cases in which there is an explanatory gap between certain facts even though it is inconceivable that the explanans obtains without the explanandum. A view like this does seem to be possible even if the examples of it I gave are not. That is all that is needed to show that (L1) is false and, therefore, that Levine's second objection to the zombie argument is unsuccessful.

8. CONCLUSION

We started with the zombie argument and Levine's no-metaphysics-from-epistemology objection. We rejected the objection but what is wrong with it is interesting. It shows that a successful objection has to allow that conceivability is evidence for possibility and also rebut Kripke's version of the zombie argument.

Levine's second objection succeeded on both counts but it contains a false premise. The premise is false for an interesting reason, though: there could be an

explanatory gap even though it is inconceivable the explanans obtains without the explanandum. This was shown by considering an aesthetic theory and metaethic. These theories are in some ways like Levine's. Consider the metaethic. Like Levine, the Moorean thinks there are objectionable explanatory gaps. Levine's view—that there is an objectionable explanatory gap between physical (or functional) facts and sensational facts—is pretty common. The Moorean view is not; it is thought to be pretty weird. But what is weird about these views is the same: they claim that certain facts imply other facts and insist that nevertheless there is an explanatory gap between them. (I think it is, in part, because of the second conjunct that the Moorean insists that moral facts are non-natural.) One upshot of the paper is that that the Moorean view is better off than people think or, alternatively, that Levine's view is worse off.

Another is that, for all Levine says, the zombie argument is sound even though it seems that something must be wrong with it. But, also, even though Levine's second objection does not work, it seems something must be right about it. There are some propositions that you believe are impossible without understanding why they are impossible. Some of these seem to be conceivable. No doubt some of these beliefs about impossibility are true and so (C2) is false. But which are the conceivable impossibilities? Why them?

CHAPTER THREE IMAGINATIVE RESISTANCE

Take a novelization of *Macbeth* and a novel, *Macbeth**, just like it except that in *Macbeth** you read that the killing of Duncan is morally permissible. Imagining what happens in *Macbeth* is easy. There is no trouble, for example, imagining that Duncan is killed or that he is killed for the basest reasons. Imagining that the killing is morally permissible, though, is something else. We resist imagining that. In short, *Macbeth**, but not *Macbeth*, induces *imaginative resistance*. Why?

That there is imaginative resistance to novels is interesting. *Prima facie*, an author has complete control in directing the reader's imagination. He can have us imagine that animals talk, that snowmen sing, that a man wakes up as a cockroach. Because there is imaginative resistance (for short, resistance), the appearance is misleading; there are limits to authorial control.

Why there is resistance is also interesting. It illuminates the natures of imagination and possibility and also of the relation between them.²²

1. WHAT RESISTANCE IS

I focus on propositional imagining—imagining that a man is killed, imagining that the killing is impermissible, and so on. Some propositions induce resistance. We resist imagining that such-and-such.

“Imagining”? I can say a lot about what I do *not* mean. Imagining is not supposing. It is no harder to suppose that the killing of Duncan is morally permissible

²² Imaginative resistance is discussed in, among other places, Currie (2002), Gendler (2000), Hume (1985a), Matravers (2003), Moran (1994) (which uses the *Macbeth* example), Tanner (1994), Walton (1994) Weatherson (forthcoming), and Yablo (2002). “Imaginative resistance” comes from Moran.

than to suppose that it is morally impermissible. It is easy enough to suppose that it is morally permissible for the purpose of a *reductio*.

Imagining can be picturing but it need not be. There is a perfectly good kind of imagining such that I can imagine that the gross national product is rising even though I cannot picture this. That is the kind of imagining I have in mind. If imagining were just picturing, there would be no contrast between imagining that the killing of Duncan is impermissible and that it is permissible. It is as hard to picture the first as the second, but only the second induces resistance.

Imagining is not believing that something is true in a fiction. What you imagine need not be part of any fiction. That this table is a bit to the left is imaginable, but the proposition that this table is a bit to the left is part of no fiction. Does believing that something is true in a fiction involve imagining it? More on that later.

Imagining is not having something seem possible. Lots of propositions seem possible to me now but I am not imagining anything. Does imagining at least involve a proposition's seeming possible? Can I imagine something that does not seem possible? More on that later.

Imagining is what I engage in when I read *Macbeth* and imagine that Duncan is killed.²³ Resistance is what happens when I read *Macbeth** and try to imagine that Duncan's killing is morally permissible. What resistance amounts to, though, is disputed.

²³ Philosophical accounts of imagination are extremely hard to come by. Walton (1990) is largely about the imagination but there Walton writes, "What is it to imagine? We have examined a number of dimensions along which imaginings can vary; shouldn't we now spell out what they have in common? Yes, if we can. But I can't. Fortunately, an intuitive understanding of what it is to imagine...is sufficient to proceed with our investigation." (19)

Hume says it is difficulty in imagining. It is *hard* to imagine that Duncan's killing is morally permissible; "a very violent effort" is required to do so. (Hume (1985a): 247)

Tamar Szabó Gendler, by contrast, says that resistance is refusal to imagine. I *refuse* to imagine that Duncan's killing is morally permissible. "The primary source of...resistance," she writes, "is not our *inability* to imagine morally deviant situations but our *unwillingness* to do so." (Gendler (2000): 56)

Kendall Walton, by further contrast, says that resistance is inability to imagine. I *cannot* imagine that Duncan's killing is morally permissible. "I might imagine [believing the killing is right]," he writes, "I cannot imagine [that it is]." (Walton (1994): 49)

The puzzle of imaginative resistance (for short, the puzzle) is explaining why resistance occurs.²⁴ Hume, Gendler, and Walton differ on what needs explaining, as between:

- (1) Certain imaginers have difficulty imagining certain propositions.
- (2) Certain imaginers refuse to imagine certain propositions.
- (3) Certain imaginers are unable to imagine certain propositions.

Hume is right that some propositions are difficult for some people to imagine. A caveman might have trouble imagining that someone sails west and winds up in the east. It does not occur to him to imagine that the world is round. However, I do not find it plausible that resistance to *Macbeth** is a matter of difficulty. It is not that with a violent effort I can imagine that the killing of Duncan is morally permissible—I am trying *very* hard, I am being very imaginative. Still, I resist it; so I suggest putting Hume's puzzle aside and focusing just on Gendler's and Walton's.

²⁴ "The puzzle of imaginative resistance" comes from Gendler.

Gendler is right that some imaginers refuse to imagine some propositions. Someone who has bet his life savings that he won't imagine that he has three hands will make every effort not to. Asked, he will refuse.

Of course, resistance to *Macbeth** does not derive from a bet. Why would I refuse to imagine that Duncan's killing is morally permissible? Gendler suggests:

GENDLER'S PROPOSAL—We resist imagining all and only those propositions such that imagining them (a) requires us to think of the actual world as somehow other than we think it is and (b) the world's being that way repulses us.²⁵

Gendler's proposal explains resistance to *Macbeth**. If it is morally permissible in any world to kill basely then, moral principles being necessary, it is actually morally permissible. But that is repulsive. So I refuse to imagine *Macbeth**.

Gendler's proposal predicts, correctly, that I do not resist *Macbeth*. Imagining that does not require me to think of the actual world as other than I think it is. Gendler's proposal predicts, correctly, that I do not resist imagining that a man wakes up as a cockroach; there is nothing repulsive about that. (Well, nothing *morally* repulsive about it.)

Two questions: first, is

(2) Certain imaginers refuse to imagine certain propositions

a good characterization of resistance? Second, *assuming it is*, does Gendler's proposal explain resistance? The answer to both questions is "no"—an example brings this out.

Consider this joke:

A man walks into a bar.

²⁵ This considerably simplifies the proposal in Gendler (2000) but keeps its main idea.

The joke—I don't know what else to call it—is not funny, and I resist imagining that it is. Better, I resist imagining that the joke, just as it is, is funny and the facts that would make it funny—facts about men and bars and bar jokes—are as they actually are.²⁶ Gendler's proposal does not predict resistance here. I am not repulsed by the joke's being funny. What kind of comic prig would be so repulsed?

It follows that Gendler's proposal is flawed. The proposal explains some moral resistance, but there is more resistance than moral resistance. Comic resistance is an example of resistance to an *evaluative* but non-moral proposition. We can mark off the evaluative/non-evaluative distinction well enough by example.

That a joke makes Alan laugh is non-evaluative; that a joke is funny is. That there is an inch of paint on the canvas is non-evaluative; that the painting is beautiful is. That a pen leaks is non-evaluative; that it is a bad pen is.

Hume thinks all resistance is evaluative resistance (Hume (1985a): 246-247), but that is a mistake. Consider:

Shape—Alan ran his hands over the three-sided square.

Shape induces resistance, but does not express an evaluative proposition. Gendler's proposal cannot explain resistance here. There is nothing repulsive about *Shape*.

²⁶ Some people, Gendler included, have no trouble imagining that the joke is funny (Gendler (2000): 79, fn. 3). My guess is that such people think that funniness is a dispositional property: for something to be funny is just for it to be disposed to make people laugh. And it is easy enough to imagine that people laugh at the joke. Insofar as one resists imagining that a joke is funny but does not resist imagining that people laugh at it, one has grounds for denying the dispositional analysis.

It is easy to imagine that a community laughs at bad jokes. Small children form such a community. And crazy people:

When [mad people] laugh at things the rest of us don't think are funny, like the death of a parent, they're not being penetrating...They're laughing because they're mad, too mad to be able to tell what's funny any more. The rewards of being sane may not be very many but knowing what's funny is one of them. (Amis (1984))

It seems there could be communities who laugh at jokes that are not funny. And it is no use saying that what is funny is just what normal people are disposed to laugh at. Having a bad sense of humor does not make you abnormal. This is a problem for the dispositional analysis of funniness.

Perhaps resistance to *Macbeth** is not akin to comic resistance or resistance to *Shape*. Perhaps—but what is the reason to think not? Introspectively, my reactions to *Macbeth**, the proposition that the joke is funny, and *Shape* are the same. The *prima facie* case is that evaluative resistance of the moral kind, evaluative resistance of the non-moral kind, and non-evaluative resistance are of a piece. Absent some reason to doubt the case, it is a problem for Gendler’s proposal that it cannot explain some non-evaluative resistance or some evaluative resistance.

Two points have emerged. First, the counterexamples to Gendler’s proposal suggest a constraint on a solution to the puzzle:

EVALUATIVE AND NON-EVALUATIVE—Some evaluative, non-moral propositions induce resistance. Some non-evaluative propositions induce resistance.

Gendler’s proposal fits badly with this constraint.²⁷

Second, resistance to *Shape* shows that resistance is not refusal. It is not that I refuse to imagine that something merely three-sided is square. Why would I refuse? I seem *unable* to imagine *Shape*. Likewise, I think, I am unable to imagine *Macbeth** and that the bar joke is funny. So

(3) Certain propositions are unimaginable for certain imaginers
is the best characterization of imaginative resistance. Our puzzle is to explain why (3) is true.

²⁷ Gregory Currie’s proposal in Currie (2002) has the same strengths as Gendler’s but, like Gendler’s, it does not fit well with EVALUATIVE AND NON-EVALUATIVE.

2. SOME IMPOSSIBILITY PROPOSALS

We started by resisting *Macbeth**. Considering Hume's and Gendler's theories of what resistance is and Gendler's proposal for why we resist, we were led to the view that resistance is inability to imagine a proposition and that we resist far more than *Macbeth**. In particular, there is more to resistance than evaluative resistance. An obvious proposal now suggests itself:

THE OBVIOUS PROPOSAL—We resist all and only the impossible propositions. We resist them because whatever is imaginable is possible, whatever impossible unimaginable.

The obvious proposal fits with EVALUATIVE AND NON-EVALUATIVE. Those evaluative, non-moral propositions that are impossible are resisted, and so are those non-evaluative, impossible propositions. For example, *Macbeth**, *Shape*, and the proposition that the bar joke is funny are impossible and so resisted.

The obvious proposal is simple, and makes use of a thesis about the relation between imaginability and possibility that many people accept anyway. That is all to the good, but the proposal has some serious flaws; we can bring them out as follows.

Alan thinks that base killing is morally permissible. *Macbeth** seems to induce no resistance in him. What seems to induce resistance is *Bad Macbeth*, a novel just like *Macbeth** except that, in *Bad Macbeth*, one reads that the killing of Duncan is morally forbidden (in *Macbeth* itself, the killing is not explicitly morally evaluated). The obvious proposal has to deny the appearances. For the propositions *Macbeth** expresses are impossible, the propositions *Bad Macbeth* expresses are not. *Macbeth** seems imaginable to Alan; according to the obvious proposal, it isn't. Alan seems to resist *Bad Macbeth*; according to the obvious proposal, he doesn't.

These consequences are implausible. Our understanding of resistance comes from our reaction to cases, including our reaction to *Macbeth**. But if it turns out that Alan, not us, is right about the permissibility of baseless killing, the obvious proposal entails that we did not resist *Macbeth**. This seems obviously, but instructively, wrong.

The Alan case shows that

POSSIBLE—Propositions that are possible can be resistance-inducers.

Alan resists the possible—indeed, necessary—proposition that baseless killing is impermissible; he resists because his moral beliefs are so skewed. What this also brings out is that resistance can vary from imaginer to imaginer, depending on the imaginer's beliefs.

RELATIVITY—Resistance is imaginer-relative: what induces resistance in one imaginer need not induce resistance in another.²⁸

The Obvious Proposal is inconsistent with both POSSIBLE and RELATIVITY. If it is true, only impossible propositions are resisted and all imaginers resist the same propositions: the impossible ones. It is easy enough to modify the obvious proposal to respect these constraints:

THE MODIFIED PROPOSAL—We resist all and only those propositions we believe are impossible. We resist them because we believe they are impossible.

Because there are a few varieties of impossibility, there are a few varieties of the modified proposal. Maybe we are unable to imagine all and only those propositions we believe are *naturally* impossible, that is, impossible given the actual laws of nature. If so, the modified proposal is better put as:

²⁸ Should resistance be time-relativized or relative to a set of beliefs? After all, Alan, that very imaginer, is able to imagine that baseless killing is wrong—he would just have to change his beliefs. I ignore this complication.

THE NATURAL IMPOSSIBILITY PROPOSAL—We resist all and only those propositions we believe are naturally impossible. We resist them because we believe they are naturally impossible.

This is a non-starter. The natural impossibility proposal clearly does not provide a sufficient condition for imaginative resistance. Science fiction stories express one natural impossibility after another, but readers, even those who know the propositions are naturally impossible, have no trouble imagining them.

Perhaps what is unimaginable is what we believe to be conceptually impossible, that is, what is *a priori* impossible. If so, the modified proposal is better put as:

THE CONCEPTUAL IMPOSSIBILITY PROPOSAL—We resist all and only those propositions we believe are conceptually impossible. We resist them because we believe they are conceptually impossible.

The proposal does not produce a necessary condition for resistance. Plenty of people do not think that moral necessities are *a priori* necessities. It is not, they say, conceptually impossible that baseless killing is right. But such people resist imagining that base killing is right. Moran, for example, does.²⁹

Perhaps what is unimaginable is what is believed metaphysically impossible, that is, false in all possible worlds. If so, the modified proposal is better put as:

THE METAPHYSICAL IMPOSSIBILITY PROPOSAL—We resist all and only those propositions we believe are metaphysically impossible. We resist them because we believe they are metaphysically impossible.

This is insufficient. I believe it is metaphysically impossible that a man wakes up as a cockroach, but that is imaginable. My reaction to *The Metamorphosis* is not at all like my reaction to *Macbeth**.

²⁹ Cf. Moran (1994): 101-102.

Gendler claims that the metaphysical impossibility proposal does not even produce a necessary condition for imaginative resistance. She claims that the following story induces resistance.

The Mice

Once upon a time there were a bunch of mice. The mice who had white fur were hardworking and industrious, but the mice who had black fur were slothful and shiftless...So it was not surprising that the mice with white fur tended to be much better off than the mice with black fur...

Even so, the mice with white fur were very generous to the mice with black fur. They gave bits of cheese to the black mouse babies. They left piles of nuts and seeds in the black mouse neighborhoods. And obviously, they provided the black mice with role models of diligence and industry. But the mice with black fur just kept to their old ways. They seemed constitutionally incapable of changing...Some of the white mice kept providing the black mice with food and other necessities, but most did not. And that was the right thing to do. For the distribution of resources in the mouse world reflected the relative merits of the two mouse groups. All the mice got what they deserved. (Gendler (2000): 73-74)

Gendler writes, "it is virtually impossible for us to take *The Mice* as anything but an extremely crude allegory for race relations. As such, the story evokes...imaginative resistance" (*ibid.*). If so, the metaphysical impossibility proposal does not produce a necessary condition for resistance—no one thinks the propositions expressed by *The Mice* are metaphysically impossible (nevermind that mice can't be shiftless). But I do not think the story induces resistance. I agree with Gendler that the story is a crude allegory for race relations. It is hard to *believe* that the story is an apt allegory. But I do not resist the story. In fact, it is because I can imagine *The Mice* that I can see it is such a crude allegory.

3. WALTON AND WEATHERSON'S PROPOSAL

The metaphysical impossibility proposal seems to give a necessary condition for resistance. Gendler's mice story is the only proposed counterexample to it, and it does not succeed. Nonetheless, the metaphysical impossibility proposal is false; its proposed sufficient condition is not a sufficient condition.

The idea that resistance has something to do with what one believes to be impossible is attractive. The proposition that base killing is morally permissible induces resistance in me. So does the proposition expressed by *Shape*. So does the proposition that the bar joke is funny. By contrast, neither *Macbeth* nor *Bad Macbeth* nor *The Mice* induce resistance. The fact that the first group of propositions is (believed to be) impossible and the second group (believed to be) possible seems important in explaining resistance. The fact that there are no convincing examples of resistance-inducers that are believed to be possible is important.

Also, the fact that an action is morally impermissible *because* it is a base killing seems important for explaining our resistance to *Macbeth**. Consider our resistance to *Shape*. The fact that an object is *not* a square *because* it has only three sides seems important for explaining our resistance to *Shape*. Consider our resistance to the proposition that the joke is funny. The fact that the joke is not funny *because* it is so lame seems important for explaining our resistance here.

Walton's proposal builds around these facts; the proposal has lately been improved by Brian Weatherston. The Walton/Weatherston (henceforth, W/W) proposal is complicated; we will need some background to understand it.³⁰

³⁰ See Walton (1994) and Weatherston (forthcoming).

Of the impossible propositions, some are *compound*. A proposition is compound iff it is logically analyzable into conjuncts.³¹ A proposition is prime otherwise. So, that Alan is tall and Alan is not tall is a compound impossibility. That something is non-identical with itself is a prime impossibility.

Among the compound impossibilities, some are asymmetric, the rest symmetric. A compound impossibility is asymmetric iff one conjunct is (or would be) false in virtue of the other being true but not vice versa. That act A is a base, morally permissible killing is an asymmetric compound impossibility. It is compound, according to W/W, being analyzable into the proposition that act A is a base killing and act A is morally permissible. It is also asymmetric: the second conjunct would be false in virtue of the first being true (an action is morally impermissible in virtue of being a base killing) but not vice versa (an action is not a base killing in virtue of being morally impermissible, but rather in virtue of the behavior of its agent).

The compound impossibilities that are not asymmetric are symmetric. For example, that something is not identical with itself and twice two is five is symmetric. Neither conjunct is false in virtue of the other being true. That something is good and not good is also symmetric.

Setting aside the subtle question of how exactly to understand “in virtue of,” the W/W proposal is:

THE W/W PROPOSAL—We are unable to imagine all and only those propositions we believe are asymmetric compound impossibilities. We are unable to imagine them because we believe they are asymmetric compound impossibilities.³²

³¹ As Weatherson notes, he needs to assume that propositions have structure and, in particular, that they can be conjunctive. This assumption is controversial. See Weatherson (2004): 4.

³² Weatherson’s proposal is importantly different. He believes that we are unable to imagine all and only asymmetric compound impossibilities. Whether or not we believe a proposition is an asymmetric

The proposal has a lot going for it. It is consistent with EVALUATIVE AND NON-EVALUATIVE, POSSIBLE, and RELATIVITY. It gets the cases of resistance and non-resistance that we have considered right.

However, it gives neither a necessary nor a sufficient condition for resistance. Not necessary: I believe that the proposition that something is not identical with itself is a prime impossibility. I believe it is not decomposable into conjuncts. Still, I resist it. That something is good and not good is, I believe, a symmetric compound impossibility. Still, I resist it.

Not sufficient: it is imaginable that there is a physical duplicate of me and it has no consciousness. I think this is an asymmetric compound impossibility. It is impossible: any physical duplicate of me is a mental duplicate and, since I am conscious, so is my duplicate. It is an asymmetric compound: the right conjunct is false in virtue of the left conjunct being true but not vice versa. If something is not a physical duplicate of me, this is in virtue of a physical difference, not a mental difference. So if it is imaginable that there is a physical duplicate of me and it has no consciousness—and I think this is—then the W/W proposal does not give a sufficient condition for resistance.

The case is controversial. Some deny that it is imaginable and others deny that it is impossible. So I do not want to rest much argumentative weight on it. Still, I think it is at least a *prima facie* problem for the W/W proposal. Our reaction to the case is *prima facie* different from our reaction to *Macbeth**.

I conclude that the W/W proposal is not right. None of the impossibility proposals we have considered is right. Some predict resistance where there is none;

compound impossibility is irrelevant to our ability to imagine it. I reject this proposal for the same reasons I reject the obvious proposal: it is inconsistent with POSSIBLE and RELATIVITY.

others fail to predict resistance where there is some. Some do both. On the one hand, an inductive argument against impossibility proposals is starting to look attractive. On the other, since we have seen no resistance-inducer that is believed to be contingent, I think that some sort of impossibility proposal is right.

4. THE ESSENTIAL PROPOSAL

Part of what is resistance-inducing about the proposition that Alan ran his hands over the three-sided square is that having just three sides is inconsistent with what it is to be a square. *Shape*, then, expresses an *essentially impossible* proposition: it represents a square as lacking one of its essential properties. Generally, a proposition is essentially impossible iff it represents something as lacking one of its essential properties; otherwise, it is essentially possible. That a square has just three sides is essentially impossible, that a square is red is essentially possible.

A property is essential to something iff it is one of the group of properties that the thing has to have to be what it is. Being four-sided, for example, is essential to being a square, as is having four angles, and having sides of equal length.

Not all metaphysical impossibilities are essentially impossible. Some properties a thing could not have lacked are not essential to it. For example, it is essential to me to be a person, that is part of what I am. It is no part of what I am that I am the sole member of singleton Tyler. However, it is impossible for me to exist without being a member of singleton Tyler. In any world in which I exist, I am a member of that set. That I exist but

am not a member of singleton Tyler, then, is essentially possible but metaphysically impossible.³³

Any essential impossibility is, I believe, metaphysically impossible. It is essentially impossible that a square is just three sided or that it has an acute angle or that it has unequal sides. All of these propositions are metaphysically impossible, too.

All resistance-inducers are believed to be essentially impossible. Take *Shape*. What it is to be square is, among other things, to have four sides. So *Shape* expresses an essential impossibility.

Suggestion: we are unable to imagine all and only those propositions we believe are essentially impossible.

This is not quite right. I believe it is essential to being a person that you have a human body but the start of *The Metamorphosis* induces no resistance in me. I can imagine that Gregor Samsa, that person, woke up with a cockroach's body. So the suggestion does not provide a sufficient condition for resistance.³⁴

The Metamorphosis is instructive here. By the end of it, there is no imagining that Samsa survives. This is not simply because he is missing essential properties. He is missing essential properties from the beginning, and the beginning induces no resistance. Rather, by the end, he is too little like Samsa, too much like a roach to imagine that it is Samsa scuttling around his room. What this suggests is that what induces resistance is not just a proposition that you believe gets the essence of something wrong but one that

³³ The singleton example comes from Fine (1994), in which the distinction between essential properties and properties a thing could not have lacked is illuminated.

³⁴ Joseph Almog seems to endorse the suggestion (or the thesis that we are unable to imagine those propositions that *are* essentially impossible) in his work on Descartes and the mind-body problem, Almog (2002a, 2002b). The Samsa case shows that the suggestion is mistaken and there are other counterexamples to it. Almog, like many people, thinks it is essential to being a person that you come from sperm and egg. So the suggestion predicts that he resists that God made Eve—a woman—from Adam's rib. Or he resists that she is a woman. I would be surprised if he resists either.

does so by confusing the thing with something else. For example, what is resistance-inducing is not only that (you believe) it is essentially impossible that Samsa lack a human body and mind. It is also that (you believe) something with all the properties the protagonist has at the end of the story is a roach.

Shape is a resistance-inducer. What is resistance-inducing is not only that (you believe) it is essentially impossible that a square lack four sides. It is also that (you believe) a three-sided figure is a triangle. Presented with *Shape*, you are confused about what you are supposed to imagine: a three-sided figure? Or a square?

That the bar joke is funny is a resistance-inducer. What is resistance-inducing is not only that (you believe) it is essentially impossible that a funny joke lack cleverness or the element of surprise or absurdity or whatever. It is also that (you believe) a joke like that is lame and stupid. Presented with the proposition that the bar joke is funny, you are confused about what you are supposed to imagine: the lame, stupid bar joke? Or a funny joke?

*Macbeth** is a resistance-inducer. In particular, that the baseless killing of Duncan is morally permissible is a resistance-inducer. What is resistance-inducing is not only that (you believe) it is essentially impossible that a morally permissible killing be done for no good reason. It is also that (you believe) a baseless killing is morally *impermissible*. Presented with the *Macbeth**, you are confused about what you are supposed to imagine: a baseless killing? Or a morally permissible killing?

The Samsa case shows that we resist only some essentially impossible propositions. The Samsa, joke, *Shape*, and *Macbeth** cases show that resistance-inducers are those that get the essence of things wrong in a certain way. They portray the subject

as something other than what it is. This subset of the essentially impossible propositions is the *essentially confused* propositions. This suggests:

THE ESSENTIAL PROPOSAL—We resist all and only those propositions we believe are essentially confused. We resist them because we believe they are essentially confused.

The Samsa case shows that the essential proposal can explain non-evaluative resistance. The joke case shows that it can explain evaluative, non-moral resistance. So the essential proposal is consistent with EVALUATIVE AND NON-EVALUATIVE.

The essential proposal is also consistent with POSSIBLE and RELATIVITY. So long as your beliefs about essence are skewed enough, the essential proposal predicts you can resist a possible proposition. You mistakenly think it is essentially confused. Such a case also shows that the essential proposal is consistent with RELATIVITY. What induces resistance can vary from imaginer to imaginer.

Unlike the W/W proposal, the essential proposal predicts the possibility of someone (a) believing that anything that is a physical duplicate is a mental duplicate while (b) being able to imagine there is a physical duplicate of herself that is not conscious. She just has to hold that it is not essentially impossible that a physical duplicate lack consciousness. She has to hold that it is not essential to a physical state to be a mental state. She can hold this even while conceding that there is a necessary connection between the two, just as there is a necessary connection between my existing and my singleton existing. So if she thinks it is essentially possible that something is in certain physical states without being conscious, she will be able to imagine that there is a physical duplicate of her that it not conscious.³⁵

³⁵ What is harder for the essential proposal is explaining why some people cannot imagine that there is a physical duplicate that is not a mental duplicate. The essential proposal predicts resistance if the person

5. OBJECTIONS TO THE ESSENTIAL PROPOSAL

The essential proposal seems subject to a counterexample. Consider this story, adapted from one by Graham Priest.

Shopping in a flea market in Indonesia, Sylvan came across a small, empty box...with an elephant figurine in one corner.

If the proposition this story expresses is imaginable, the essential proposal is false. What it is to be an empty box is to be a box with nothing in it and what it is to be a non-empty box is to be a box with something in it. Everyone believes this and so everyone believes the proposition expressed is essentially confused. There is a simple argument that the proposition is imaginable: it is true in the story that there is an empty box with something in it. So it is imaginable that there is an empty box with something in it. So the essential proposal is false. Gendler argues like this against the conceptual impossibility proposal (Gendler (2000): 68).

I agree that it is true in the story that there is an empty box with something in it. I just deny that it follows that it is imaginable that there is an empty box with something in it. According to Gendler, any proposition accepted as true in a fiction is imaginable. I deny this. Take someone who cannot imagine that there are physical duplicates of people but without consciousness. Now consider this story:

Albert was a physical duplicate of Alan but without consciousness. He spent most days in front of a TV.

believes that it is essential to being a physical duplicate that you are a mental duplicate and also that what it is to be a mental non-duplicate is to be a physical non-duplicate. If so, the person resists imagining that there is a physical duplicate that is not a mental duplicate. The proposition is resistance-inducing not only because (he believes) it is essentially impossible that a physical duplicate is not a mental duplicate. It is also because (he believes) a mental non-duplicate is a physical non-duplicate. He thinks, "What is it I am supposed to imagine? A physical duplicate? Or a physical non-duplicate?"

I assume that our imaginer can't imagine the proposition the first sentence expresses. Still, can't she see what is true in the story? Asked, "What happened in the story?" she won't say, "I have no idea" or "Someone was in front of a TV but I can't tell you anything about him." She will say, "There was a physical duplicate of Alan who had no consciousness and spent most days in front of a TV. (What a story.)" So what is accepted as true in a fiction need not be imaginable. So the apparent counterexample is no counterexample to the essential proposal.

A second objection to the essential proposal is as follows. Some people do not believe in essences. So they do not believe any proposition is essentially confused. If these people experience resistance, then the essential proposal is false. In particular, it does not give a necessary condition on imaginative resistance.

Analogous of this objection can be put to any impossibility proposal (except the obvious proposal). There are people who do not think any proposition is impossible: they have no truck with modality (or propositions). There are others who think that ethical claims do not have truth values and, hence, think that there are no impossible ethical propositions. So if they resist *Macbeth**, it is not because they believe it expresses an essentially confused proposition; they believe the sentence, "The killing of Duncan was morally permissible" expresses no proposition. If these people experience resistance, so much the worse for impossibility proposals—they do not give necessary conditions for resistance.

It might be that this objection is a good one (the case that some impossibility proposal is right is pretty strong, though). If so, all the impossibility proposals, except the obvious proposal, are false. If so, the essential proposal is the best of a bad bunch.

At worst, it gives only a sufficient condition for resistance. Since most people believe in essences, the essential proposal will be able to explain resistance in most people. Of the impossibility proposals, the essential proposal is best. It gets all the cases we have considered right.

6. CONCLUSIONS

I began by clarifying the nature of imaginative resistance. I then criticized Gendler's proposal, the obvious proposal, and some other impossibility proposals. Each was rejected but what is wrong with each is interesting and leads to some constraints on a solution to the puzzle.

I suggested a solution. If it is right, then a thing's essential properties are not simply properties it has to have and there is a sort of impossibility, essential impossibility, that is neither conceptual nor metaphysical nor natural.

Also, if the essential proposal is right, imagination reveals beliefs not so much about what is possible but about what is essential to things. I can imagine that I exist without singleton Tyler. This is not because I think this is possible—I know it isn't—but because I believe what I am has nothing to do with sets. By contrast, I cannot imagine that singleton Tyler exists without me. So I have reason to think it is essential to singleton Tyler to have me as a member.

I can imagine that I am in just the brain state I am in when I have a headache but without feeling any pain at all. This is not because I think this is possible—I don't think it is—but because I believe that what pain is has nothing to do with brains. By contrast, I cannot imagine that the brain state exists without being a state of a brain. So I have

reason to think it is essential to the state to be instantiated by a brain. I think this is essentially Kripke's point at the end of *Naming and Necessity*: what it is to be in pain differs from what it is to be in a certain brain state. Different essences, so different states. It does *not* follow that there can be pain without a certain brain state. From their different essences, it does not follow that I can exist without singleton Tyler or singleton Tyler without me.

Regardless of whether the essential proposal is right, resistance itself illuminates the nature of the imagination. It shows that imagining is distinct from accepting as true in a fiction, distinct from seeming possible, and from supposing.

Distinguishing imagining from supposing, for example, is necessary for preserving the imaginability/possibility link; there is no supposability/possibility link. Making the distinction between imagining and supposing, however, raises the question of what the difference comes to and how it allows for an imaginability/possibility link.

A case of moral imagination we have considered seems to me to show that the link is not entailment. Alan the base killer can imagine that baseless killing is morally permissible. So at least one impossibility is imaginable. Friends of the thesis that imaginability entails possibility will respond in one of four ways.

First way: Alan is in some sense imagining an impossible proposition. However, only a certain sort of imaginability entails possibility and Alan is not engaging in that sort.

Second way: Alan is not *imagining* what he seems to be imagining. Instead, he is in some mental state like imagining but one that does not entail possibility.

Third way: Alan is confused about what he is imagining. He is imagining a possible proposition, not the impossibility he seems to be imagining.

Fourth way: Alan is imagining what he thinks he is imagining and, appearances to the contrary, the proposition is possible.

One or more of these responses might be plausible in cases of non-moral imagining, for example, in response to a case in which someone seems to imagine that water is not H₂O or that Queen Elizabeth is the daughter of the Trumans. In the moral case, the fourth, at least, is out of the question. If any proposition is impossible, that a baseless killing is permissible is.³⁶

The third reply is not much more plausible. It is hard to see just what possible proposition Alan is confusing with that a baseless killing is permissible. Maybe he is confusing that something just like a baseless killing is permissible with that a baseless killing is permissible but so what? If the second is impossible, so is the first.

Mightn't Alan be confusing imagining that a baseless killing is permissible with imagining that he believes that a baseless killing is permissible? No—there is a big difference between those propositions and Alan is bright enough and alert enough to see the difference between the two propositions.

³⁶ Objection: "Not so: the case supports moral relativism: that a baseless killing is permissible *is* possible *relative to Alan's moral standards*." Putting aside for a moment whether this is a good reply, if relativism is a consequence of the thesis that conceivability entails possibility that is pretty interesting.

However, it is not a good reply. At least, adopting moral relativism only gives the thesis temporary relief. Consider a non-moral case. Alan believes there cannot be composite objects. He seems to be able to imagine that there are exactly two objects, two simple particles. Bob the mereological universalist seems to be able to imagine that there are exactly three objects, two simple particles and the thing they compose. It is not at all plausible that that there are exactly two objects is possible *relative to Alan's beliefs* and impossible *relative to Bob's beliefs*. The proposition is either impossible or it isn't. So either the second, third, or fourth reply has to be invoked by someone who believes that conceivability entails possibility. For reasons analogous to those given in the moral case, I think each is implausible.

We are left with the first and second replies. The second is less plausible. If it is right, we have to explain away Alan's seeming to imagine that a base killing is permissible as Alan's bearing some mental state other than imagining towards that proposition. Which? And how did bright, alert Alan confuse it with imagining? The problem here is that the mental state has to be enough like imagining so that bright, alert Alan plausibly confuses it with imagining but different enough from imagining so that it does not entail possibility.

We are left with the first reply. According to it, Alan is imagining that a base killing is permissible but not, say, clearly and distinctly imagining that. Only clear and distinct imagining entails possibility. (Maybe only clear and distinct imagining is even a guide to possibility. I doubt this second claim. We do take imagining the content of novels to be a guide to possibility—when you read novels, you sometimes learn how things could be—and the sort of imagining Alan engages in is precisely the sort that imagining the content of novels involves.)

What I don't like about this reply is that I don't see why Alan's imagining need be any less clear and distinct than my imagining that a baseless killing is impermissible. The difference between us is just that Alan has false beliefs about baseless killing. Why should that impact the clarity and distinctness of his imagination? *Prima facie*, then, Alan's moral imagination shows that imaginability does not entail possibility.

Say that the appearance is veridical: imaginability does not entail possibility. This leads to a general issue. Some impossibilities are imaginable, some aren't. What is the difference? The essential proposal provides an answer.

REFERENCES

- Almog, Joseph. 2002a. *What Am I?* Oxford: Oxford University Press.
- Almog, Joseph. 2002b. "Pains and Brains," *Philosophical Topics* 30: 1-29.
- Amis, Kingsley. 1988. *Stanley and the Women*. New York: Harper Collins.
- Blackburn, Simon. 1973. "Moral Realism," reprinted in Blackburn (1993).
- Blackburn, Simon. 1984. *Spreading the Word*. Oxford: Clarendon Press.
- Blackburn, Simon. 1985. "Supervenience Revisited," reprinted in Blackburn (1993).
- Blackburn, Simon. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Block, Ned. 1994. "Consciousness" in Guttenplan (1994).
- Brink, David. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Chalmers, David. 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Chalmers, David, ed. 2002. *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- Currie, Gregory. 2002. "Desire in Imagination," in Gendler and Hawthorne (2002).
- Davies, Martin and Glyn Humphries, eds. 1993. *Consciousness*. Oxford: Blackwell.
- Fine, Kit. 1994. "Essence and Modality," *Philosophical Perspectives* 8: 1-16.
- Foot, Phillipa. 2000. *Natural Goodness*. Oxford: Oxford University Press.
- Geach, Peter. 1956. "Good and Evil," *Analysis* 17.
- Gendler, Tamar Szabo. 2000. "The Puzzle of Imaginative Resistance," *Journal of Philosophy* 97: 55-81.
- Gendler, Tamar Szabo and John Hawthorne. 2002. *Conceivability and Possibility*. Oxford: Oxford University Press.
- Guttenplan, Samuel, ed. 1994. *A Companion to the Philosophy of Mind*. Oxford: Blackwell.
- Hare, R.M. 1984. "Supervenience," *Proceedings of the Aristotelian Society Supplement* 58: 1-16.
- Horgan, Terence. 1993. "From Supervenience to Superdupervenience," *Mind* 102: 555-586.
- Horgan, Terence and Mark Timmons. 1992. "Troubles on Moral Twin Earth: Moral Queerness Revived," *Synthese* 92: 221-260.

- Hume, David. 1985a. "Of the Standard of Taste," in Hume (1985b).
- Hume, David. 1985b. *Essays Moral, Political, and Literary*. Indianapolis: Liberty Fund.
- Jackson, Frank. 1998. *From Metaphysics to Ethics*. Oxford: Oxford University Press.
- Jackson, Frank and Philip Pettit. 1995. "Moral Functionalism and Moral Motivation," *Philosophical Quarterly* 45: 20-40.
- Kieran, Matthew and Dominic Lopes. 2003. *Imagination, Philosophy, and the Arts*. London: Routledge.
- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Levine, Joseph. 1983. "Materialism and Qualia: The Explanatory Gap" in Chalmers (2002).
- Levine, Joseph. 1993. "On Leaving Out What It's Like" in Davies and Humphries (1993).
- Levine, Joseph. 2001. *Purple Haze*. Oxford: Oxford University Press.
- Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. London: Penguin.
- Mackie, J.L. 1980. *Hume's Moral Theory*. London: Routledge and Kegan Paul.
- Mackie, J.L. 1982. *The Miracle of Theism*. Oxford: Oxford University Press.
- Matravers, Derek. 2003. "Fictional Assent and the (So-Called) Puzzle of Imaginative Resistance" in Kieran and Lopes (2003).
- Moore, G.E. 1993. *Principia Ethica*, revised edition. Cambridge: Cambridge University Press.
- Moran, Richard. 1994. "The Expression of Feeling in the Imagination," *Philosophical Review* 103: 75-106.
- Nagel, Thomas. 1974. "What Is It Like To Be A Bat?" in Nagel (1979).
- Nagel, Thomas. 1979. *Mortal Questions*. Cambridge: Cambridge University Press.
- Papineau, David. 1996. "A Universe of Zombies?" *Times Literary Supplement* June 21, 1996.
- Price, Richard. 1969. *A Review of the Principal Questions in Morals*, excerpted in Raphael (1969).
- Raphael, D.D. 1969. *The British Moralists, 1650-1800*. Oxford: Clarendon.

- Sturgeon, Nicholas. 2003. "Moore on Ethical Naturalism," *Ethics* 113: 528-556.
- Tanner, Michael. 1994. "Morals in Fiction and Fictional Morality II," *Aristotelian Society Supplement* 1994: 51-66.
- Thomson, Judith Jarvis. 1997. "The Right and the Good," *Journal of Philosophy* 94: 273-296.
- Walton, Kendall. 1990. *Mimesis As Make-Believe*. Cambridge, MA: Harvard University Press.
- Walton, Kendall. 1994. "Morals in Fiction and Fictional Morality I," *Aristotelian Society Supplement* 1994: 27-50.
- Weatherson, Brian. 2004. "Morality in Fiction and Consciousness in Imagination," talk delivered at the Pacific APA.
- Weatherson, Brian. Forthcoming. "My Favorite Puzzle," *Philosopher's Imprint*.
- Yablo, Stephen. 2002. "Coulda, Woulda, Shoulda" in Gendler and Hawthorne (2002).