

The knowledge argument, the open question argument, and the moral problem

Michael Pelczar

Received: 5 December 2007 / Accepted: 15 July 2008 / Published online: 1 August 2008
© Springer Science+Business Media B.V. 2008

Abstract Someone who knew everything about the world's physical nature could, apparently, suffer from ignorance about various aspects of conscious experience. Someone who knew everything about the world's physical and mental nature could, apparently, suffer from moral ignorance. Does it follow that there are ways the world is, over and above the way it is physically or psychophysically? This paper defends a negative answer, based on a distinction between *knowing the fact that p* and *knowing that p*. This distinction is made intelligible by reference to criterial connections between the possession of moral or phenomenal knowledge, and the satisfaction of cognitively neutral conditions of desire and experiential history. The existence of such connections in the moral case makes for an efficient dissolution of the so-called moral problem.

Keywords Knowledge argument · Open question argument · Moral problem

This paper concerns a class of arguments intended to refute various forms of ontological naturalism, according to which phenomena of a disputed kind are of some purely natural (e.g., physical) character. An argument of this class contains two premises. The first states that knowing all of a given domain of natural facts is compatible with suffering from ignorance in a certain regard. The second premise states that if the first is true, then there are ways the world is over and above the way it is with respect to the given domain.

Call any argument of this form a *knowledge argument*. The clearest example of such an argument is Frank Jackson's knowledge argument against physicalism. Physi-

M. Pelczar (✉)
Department of Philosophy, National University of Singapore, Singapore, Singapore
e-mail: phimwp@nus.edu.sg

calism, as I understand it here, is the view that the phenomenal facts that obtain in our world—facts of phenomenal consciousness—are a species of physical facts. Thus, according to physicalism, any possible world physically indistinguishable from ours contains all of the conscious experiences that our world contains, having all the phenomenal properties that these experiences actually possess. Jackson argues against physicalism as follows:

Mary is a brilliant scientist who is, for whatever reason, forced to investigate the world from a black and white room via a black and white television monitor. She specializes in the neurophysiology of vision and acquires, let us suppose, all the physical information there is to obtain about what goes on when we see ripe tomatoes, or the sky, and use terms like ‘red’, ‘blue’, and so on. She discovers, for example, just which wavelength combinations from the sky stimulate the retina, and exactly how this produces via the central nervous system the contraction of the vocal chords and expulsion of air from the lungs that results in the uttering of the sentence ‘The sky is blue’... What will happen when Mary is released from her black and white room or is given a color television monitor? Will she *learn* anything or not? It seems just obvious that she will learn something about the world and our visual experience of it. But then it is inescapable that her previous knowledge was incomplete. But she had *all* the physical information. *Ergo* there is more to have than that, and Physicalism is false.¹

Call the argument Jackson offers here the *anti-physicalist argument*. It boils down to this:

- P1 Someone who knew all the physical facts could suffer from phenomenal ignorance (e.g., fail to know what it is like to experience red).
- P2 If someone who knew all the physical facts could suffer from phenomenal ignorance, then at least some phenomenal facts are not physical facts.
- P3 So, at least some phenomenal facts are not physical facts.

Another example of a knowledge argument is evoked by the following passage from Wittgenstein’s *Lecture on Ethics*, in which he puts forth a version of what has become known as the Open Question Argument:

Suppose one of you were an omniscient person and therefore knew all the movements of all the bodies in the world dead or alive and that he also knew all the states of mind of all human beings that ever lived, and suppose this man wrote all he knew in a big book, then this book would contain the whole description of the world; and what I want to say is, that this book would contain nothing that we would call an *ethical* judgment or anything that would logically imply such a judgment...If for instance in our world-book we read the description of a murder with all its details physical and psychological, the mere description of these facts will contain nothing which we could call an *ethical* proposition. The murder will be on exactly the same level as any other event, for instance the falling of a stone...Our words used as we use them in science, are vessels

¹ Jackson (1982, p. 130). Jackson no longer endorses this argument, or anti-physicalism.

capable only of containing and conveying meaning and sense, *natural* meaning and sense. Ethics, if it is anything, is supernatural...²

Here, Wittgenstein argues for a disjunctive conclusion: *either* there are no moral facts (“ethics is nothing”) *or* moral facts are non-natural (“ethics is supernatural”). His premises are (1) that someone could know all natural facts (i.e., the physical and mental facts together with all that they logically entail) without making any moral judgements; and, (2) that if there are moral facts, then, given (1), they are supernatural. Wittgenstein’s conclusion is, in effect, that there is no way to be a moral realist without denying *moral naturalism*, defined as the view that there are moral facts and that they are simply a species of natural (mental or physical) facts—i.e., the view that the ways the world is morally are ways it is naturally.³

It will be useful to put a name to the reader of Wittgenstein’s world-book, who plays a role in Wittgenstein’s argument akin to Mary’s role in Jackson’s. Call her Carrie. Carrie, then, is a hypothetical being who knows every natural fact, but suffers from moral ignorance in some respect. On the face of it, such a being does seem logically possible. It is true that we cannot put ourselves in the shoes of a being who knows all natural facts, any more than we can put ourselves in the shoes of someone who knows all physical facts, but this should not make us doubt that such knowledge could exist in a morally ignorant person. We should at least be able to agree that it is possible for someone who knows all the *relevant* natural facts about some wrong act not to know that the act is wrong. A moral nihilist, for example, can know all the relevant natural facts about some wrong act without believing, and thus without knowing, that it (or anything else) is wrong. God’s goodness is not supposed to be an automatic consequence of his natural omniscience, but a further quality of his not guaranteed by his omniscience.

It is tempting to turn Wittgenstein’s argument into an attack on moral realism, by adding a premise to the effect that the moral facts *are* natural facts, if they exist at all. But I am interested in seeing what happens if we cast the argument as a challenge to moral naturalism, understood as the view that the moral facts that obtain in our world are just a special kind of natural facts. Thus construed, the argument looks like this:

- M1 Someone who knew all the natural facts could suffer from moral ignorance.
- M2 If someone who knew all the natural facts could suffer from moral ignorance, then at least some moral facts are not natural facts.
- M3 So, at least some moral facts are not natural facts.

This is the knowledge argument against moral naturalism, or, for short, the *anti-naturalist argument*.

My goal in this paper is to show that the anti-naturalist and anti-physicalist arguments both fail, and for the same reason. The point of discussing both arguments in

² Wittgenstein (1965/1930, pp. 6–7).

³ While Wittgenstein’s argument is clearly reminiscent of, and probably at least indirectly indebted to, Moore’s famous open question argument, the connections between the two arguments are not entirely straightforward, and drawing them out would raise interpretive issues regarding Moore’s position that would distract from the main argument of this paper. For Moore’s version, see the first chapter of Moore (1903). It is interesting that C.D. Broad frames a version of the anti-physicalist knowledge argument in terms of an open question: see Broad (1925, pp. 614–615).

one place is to stress that their failure is structural, rather than tied to some specific feature of moral or phenomenal discourse or reality. The arguments follow a pattern that might be reproduced in the service of other forms of dualism (or, taking the further step at which Wittgenstein hints, anti-realism); showing that they belong to a flawed *class* of arguments will close off these blind alleys.⁴

Before developing my response to the knowledge arguments, however, I review what is perhaps the most popular style of response to the knowledge arguments. This is the so-called “old fact/new guise” strategy. In the next section, I argue that the old fact/new guise strategy fails. Clarifying the strategy will also allow me to differentiate it from the strategy I favor.⁵

In Sect. 2, I state and defend my own response to the anti-naturalist argument. This is what I call the *old relatum/new relation* strategy.

In Sect. 3, I explain how the basic idea behind the old relatum/new relation strategy dissolves the moral problem (as it has been dubbed by Michael Smith).

In Sect. 4, I apply the old relatum/new relation strategy to the anti-physicalist argument.

Finally, in Sect. 5, I clarify the differences between the old fact/new guise and old relatum/new relation strategies, and criticize some existing philosophical positions for failing to make it clear which strategy they employ.

1 Old facts, new guises

The old fact/new guise objection is probably the most popular objection to both the anti-physicalist argument and the anti-naturalist argument (or ancestors of the anti-naturalist argument). As an objection to the anti-physicalist argument, the objection is that someone who knows all physical facts—and so, given physicalism, all facts—can suffer from phenomenal ignorance simply by virtue of failing to know a truth whose truth-maker (some physical fact) is also the truth-maker of some *other* truth that she *does* know. Thus there are two senses in which one might be said to have “complete knowledge” of the physical world. In one sense, you have complete knowledge of the physical world if every physical fact is a fact you know to obtain. In another sense, having complete knowledge of the physical world requires more than knowing every physical fact: it also requires knowing every truth that the physical facts make true. Some of these truths differ from one another not with respect to what states of affairs make them true, but only with respect to the manner of their construction out of

⁴ It is possible to interpret the central argument of Quine (1960, pp. 26–79) as a knowledge argument extended to an anti-realist conclusion: (1) a radical translator can know all of a native speaker’s broadly physical (physiological, behavioral, and environmental) features without knowing what the speaker means by his utterances; thus, (2) if there is any such thing as what the speaker means by his utterances, this is something over and above his broadly physical features; but, (3) if there is such a thing as what the speaker means by his utterances, this must reduce to something about his broadly physical features; therefore, (4) there is no such thing as what the speaker means by his utterances.

⁵ Other responses to the knowledge arguments that I shall not discuss in detail are antirealism about qualia or moral value, skepticism about the claim that someone could know all the physical (or natural) facts without having all phenomenal (or moral) knowledge, and objections based on so-called ability hypotheses (e.g., Nemirow 1980; Lewis 1990).

linguistic, conceptual, or logical components or “guises.” During her imprisonment, Mary’s knowledge of the physical world was complete in the first sense, but not the second. She does learn something upon her release, but only by acquiring knowledge of truths previously unknown to her, and not by acquiring knowledge of truths with truth-makers distinct from those of the truths she already knew.⁶

The problem with this objection is that it forgets that Mary is supposed to know *all* the facts. She is supposed to know all the facts, because she is supposed to know all the physical facts, and these are supposed (by physicalism) to be all the facts there are. But if she knows *all* the facts, she knows all the facts about all the guises, all the facts about the facts they guise, and all the facts about how all these facts and guises relate to one another and to everything else. If it is a fact that **this** (looking at something red) is what it is like to experience red, she knows it. In short, she knows every fact about everything that can be known.⁷

But how on Earth could she know every fact about everything that can be known, without knowing everything that can be known? There is just no way; *knowing all the facts entails knowing all the knowables*. Mary is not in the position of an ordinary person who knows there is water without knowing there is H₂O. An ordinary person does not know all the facts. Someone who knew all the facts would know the fact that water is H₂O, is called “H₂O” by professional chemists, consists of molecules comprising two hydrogens and an oxygen, etc. By allowing that there are truths that Mary does not know, the new guise approach therefore implicitly concedes that there are *facts* that she does not know, which is all that the anti-physicalist argument needs in order to reach its conclusion.

This criticism of the old fact/new guise objection does not hinge on the controversial “new fact thesis,” according to which it is impossible to acquire new truth-apt knowledge (“knowledge-that”) without also gaining knowledge of some new fact—some fact previously unknown to the individual acquiring the knowledge. As will become clear later, I reject this thesis. What I affirm is that, whatever the objects of knowledge may be—guised facts, fine-grained propositions, epistemic intensions, or what have you—it is impossible to undergo an increase in the objects of your knowledge without adding to the list of coarse-grained, sets-of-possible-worlds-style *facts* that you know.

⁶ For the old fact/new guise objection to the anti-physicalist argument, see esp. Horgan (1984, pp. 149–152) and Lycan (1990, pp. 113–114, 117–122).

⁷ In his original paper, Jackson explicitly attributes to Mary only knowledge of all physical facts having to do with “what goes on when we see ripe tomatoes, or the sky, and use terms like ‘red,’ ‘blue,’ and so on.” (Jackson 1982, p. 130) This left it open to a physicalist to suggest that whatever ignorance Mary suffered from during her captivity was due to her not knowing certain *physical* facts not included among the facts mentioned in the preceding quote. To head off this objection, Jackson later attributes to Mary knowledge of “everything there is to know about the physical nature of the world,” so that she “knows all the physical facts about us and our environment, in a wide sense of ‘physical’ which includes everything in *completed* physics, chemistry, and neurophysiology, and all there is to know about the causal and relational facts consequent upon all this, including of course functional roles.” (Jackson 1986, p. 291) An anonymous referee points out that it is easier to get a grip on what would be involved in knowing all the neural and optical facts than it is to get a grip on what would be involved in knowing *all* the physical facts, and that this might cloud intuitions regarding P1. If so, this is another, and independent, reason not to be persuaded by the knowledge arguments. Based on my own observations, however, there are many who do not find the move from “all the neural and optical facts” to “all the physical facts” intuition-clouding; the response to the knowledge arguments I offer below is addressed to this audience.

This is enough to thwart the old fact/new guise strategy, since the whole idea behind this strategy is to argue that Mary acquires knowledge of new (fine-grained) truths but not of new facts.⁸

One might also attempt to block the anti-naturalist argument with an old fact/new guise objection, and various moral naturalists have in fact objected to Moore's Open Question argument in this way. Applied to the anti-naturalist argument we are considering, the objection goes that Carrie's moral ignorance is a consequence of her lacking moral concepts and so, in a sense, knowledge of moral truths, but does not necessarily involve any failure on her part to know a moral fact. A single moral fact can figure as the truth-evaluable substance of more than one knowable—more than one moral truth—so that there is a one-many correlation between moral facts and moral truths. Incorporated into one conceptual, linguistic, or logical structure, a moral fact may constitute one truth, while it might constitute a different truth when incorporated in a different structure. Analogously, one might know that Deep Throat supplied information to *The Washington Post* without knowing that Mark Felt did so, as a result of knowing a single fact—the fact that Deep Throat supplied information to the *Post*, a.k.a. the fact Mark Felt supplied information to the *Post*—under a guise that corresponds to the description “the secret source whose insider guidance was vital to *The Washington Post*'s prize-winning coverage of the Watergate scandal” but not under a guise that corresponds to the description, “the associate director of the Federal Bureau of Investigation who retired in 1973.” Hence we can concede that Carrie knows every fact under some guise, without being forced to the conclusion that she knows every fact under *every* guise. That is, we can concede that Carrie knows, for each fact, including each moral fact, some truth which that fact makes true, without having to concede that Carrie knows every truth. Carrie is morally ignorant because she lacks moral concepts—the concepts that constitute the guises of the moral truths that she does not know. But from this it does not follow that there is some fact unknown to her, moral or otherwise.⁹

The reply to this is essentially the same as before: if naturalism is true, all the facts are natural facts, in which case Carrie knows all the facts. But if Carrie knows all the facts, she knows all the facts about which facts appear under which guises, or are incorporated into which conceptual structures. In short, if she knows all the facts, she knows all the facts about all the truths. But, for each truth, it is a fact that it is precisely the truth that it is, entailing whatever facts or truths it entails, entailed by whatever facts or truths entail it, relating in whatever ways it does to conceptual guises, logical structures, etc. But if she knows all these facts about a truth, then, presumably, she knows that truth; what else is there to know about it? Thus, if Carrie knows all the

⁸ For the new fact thesis, see Lockwood (1989, pp. 136–137), Thau (2002, p. 127), and Chalmers (2004, p. 289).

⁹ Instead of descriptions, we could think of guises as individual concepts, indexical modes of presentation, epistemic intensions, conceptions of one's own abilities, or some logical construction out of these; see, for example, Putnam (1975, p. 280), Harman (1977, p. 19), Miller (1985, pp. 514–516), Sturgeon (1988, pp. 240–243), Lycan (1986, pp. 80–81), Boyd (1988, pp. 199, 210, 223), Brink (1989, pp. 156–167), Dreier (1992), Smith (1994, p. 38), Sayre (1997), and Brink (2001). This defense of moral naturalism arguably goes all the way back to Frankena (1939).

facts, she must know all the truths, contrary to what the old fact/new guise objection requires.

Why is the old fact/new guise strategy so popular? I suggest it is because of a tendency to slide illicitly from the plausible claim that

- (1) Every fact is such that it is possible to know it without knowing all the truths it makes true.

to the implausible claim that

- (2) It is possible to know every fact without knowing all the truths the facts make true.

Once it is recognized that these are distinct claims, however, it becomes clear that the examples routinely adduced in support of (2)—examples in which, for instance, someone knows he has water in his glass, but not that he has H₂O—really only support (1), which falls far short of what the old fact/new guise strategy requires.¹⁰

2 Old relata, new relations

Suppose, for the sake of illustration, that utilitarianism is true. The fact that recreational torture is wrong is then just the fact that recreational torture erodes (or fails to maximize) utility. But this is a purely natural fact. So, if Carrie knows all the natural facts, then she knows this fact, which is a moral fact. It follows that, given utilitarianism, every moral fact is a fact that Carrie knows. Therefore, if moral naturalism is true, Carrie *must* know all the moral facts, given that she knows all the natural facts. Given the arguments of the previous section, she must also know all the moral truths—all the knowables that have as their truth-makers relevant facts of utility. But I deny that this implies that moral naturalism precludes any moral ignorance on Carrie's part. The main idea behind my response to the anti-naturalist argument is that M2 is false, because *knowing a moral truth is necessary, but not sufficient, for having moral knowledge*. Given that to be wrong is to have natural quality F, knowing the fact that x is F is necessary, but not sufficient, for knowing that x is wrong.¹¹

Call this the *old relatum/new relation* strategy, or the “new-relation” strategy, for short. How does it differ from the old fact/new guise (henceforth: “new-guise”) strategy? The difference is in the structure of the knowledge that the strategies respectively see Carrie as lacking. On both views, Carrie lacks some truth-apt knowledge. Also on both views, every fact is a fact that Carrie knows, despite her moral ignorance. The difference between the strategies comes out if we suppose that Carrie becomes cured of her moral ignorance. What does this change involve?

The new-guise theory and the new-relation theory agree that it involves Carrie's entering into a knowledge state that she was not in before. The question over which the two theories disagree is how the new knowledge state differs from the knowledge

¹⁰ For similar criticisms of the old fact/new guise strategy, see Pelczar (2005, pp. 37–38) and Jackson (2005, pp. 318–319).

¹¹ Henceforth, I speak indifferently of “knowing all the facts” and “knowing all the truths,” since, by the argument of the previous section, these amount to the same.

states that Carrie was already in. Like any knowledge state, the new one has a four-part structure: it consists of (1) a knower (in this case Carrie), (2) a knowledge relation, (3) a fact known, and, (4) a guise of a fact.¹² According to the new-guise theory, Carrie's new knowledge state—the one she enters so as to achieve relief from moral ignorance—differs from the knowledge states she was already in only with respect to component (4). On this view, Carrie enters a new knowledge state by virtue of coming to bear the relation of knowledge to an entity to which she did not previously bear this relation. This entity may either be thought of as a complex object (having a {fact, guise} structure) of a two-place knowledge relation, or as one of two distinct objects (a fact, and a guise) of a three-place knowledge relation. Whatever the entity is, it cannot be a mere fact or state-of-affairs, since the new-guise theory allows that Carrie already knew all of these. Rather, it is a truth that has for its truth-maker some fact that was already known to Mary in the guise of some other truth.

By contrast, the new-relation theory holds that Carrie's new knowledge state does *not* differ from any of her old ones with respect to component (3) or (4). Whatever counts as an object or content of Carrie's knowledge after her moral enlightenment also counted as such before her enlightenment. Rather, the knowledge state Carrie enters so as to achieve moral enlightenment differs from any of the states she was in before with respect to component (2). She learns inasmuch as she comes to bear to some knowable a relation of knowing that she did not previously bear to that knowable, despite having all along borne to it some *other* knowledge relation. In particular, Carrie comes to bear to some purely natural fact a relation of knowledge that a person can bear to that fact only if she desires that fact not to obtain—or, so I shall now argue.

The key to making the new-relation strategy work is to identify necessary conditions of moral knowledge that are not necessary conditions of knowing all facts. Suppose we can find some condition, M, that, as a matter of logical necessity, anyone who knows that recreational torture is wrong satisfies, but that someone who knows all the facts could fail to satisfy. Then we may argue that there is no difference between (1) knowing that torture is wrong, and (2) knowing some purely natural fact while satisfying M. We could then account for Carrie's moral ignorance simply by observing that she does not satisfy M, without having to suppose that there is any fact (or truth) she does not know. At least, nothing that proponents of the anti-naturalist argument have said (or obviously could say) would prevent us from accounting for her moral ignorance in this way.

I suggest that a good candidate for condition M is the having or lacking of relevant desires. Take, for example, the desire that acts of torture not be performed. This, I shall argue, is a desire that one logically must have in order to believe, or know, that torture is wrong. But, as I shall also argue, it is also a desire that someone who knew every natural fact—including every moral fact, if there are such facts and they are natural—could lack. There is, arguably, a similar connection between believing

¹² An anonymous referee points out that this analysis allows for a perspicuous regimentation of reactions to the thought-experiments that lie behind the knowledge arguments: dualists react by positing a new fact; mainstream physicalists (or naturalists) react by positing a new guise; I react by positing a new relation (no one reacts by positing a new knower).

(and knowing) that something is right, and *not* desiring that this thing *not* be done. But I am going to focus on the case of knowing that something is wrong.

Claim 1 Necessarily, anyone who knows that X is wrong desires that X not be done.

It might be objected that a person could know that it is wrong to cheat, yet desire to cheat; more generally, you can be tempted by, and hence desire, what you know is wrong. But this objection misses the mark. The claim is not that knowing that X is wrong requires *not* desiring that X *be* done. It is that knowing that X is wrong requires *desiring* that X *not* be done. Claim 1 is consistent with someone's knowing that X is wrong, but desiring to do X; it just implies that such a person has conflicting desires: a desire to do X, and a desire that X not be done. It is true that we sometimes say things like: "Although I know it is wrong not to pay my taxes, I do not want to pay them," but, as ordinarily used, this statement is ambiguous. On one reading (the more natural one), it says that although I know it is wrong not to pay my taxes, I have a desire not to pay them at all; this is consistent with Claim 1. On another reading, the statement says that although I know it is wrong not to pay my taxes, I have no desire at all to pay them; this is inconsistent with Claim 1, but also, for the reasons given below, false.

What of the satanic individual who knows what is wrong, and whose *only* desire is to do what is wrong? It is true that Claim 1 implies that such an individual is impossible, but this is the correct verdict. The committed satanist is someone who values satisfying one sort of desire (a desire to do what he knows is wrong) at the expense of another sort of desire (a desire not to do the very things he knows are wrong). It is true that someone could be such that, for every fact of moral wrongness he knows, he desires the obtainment of that fact and has no countervailing desires. Such a person is not satanic, but morally blind. He is a monster, perhaps, but not a moral monster; a natural rather than a moral evil.

Claim 1 may be understood as a weak form of motivational internalism about moral judgement. To deny it would be to endorse an extreme form of motivational externalism, by which you could know, in a purely intellectual way, that X is wrong, yet have no desire, however feeble, unconscious, or dominated by contrary desires, that X not be done. But what would be the difference between knowing in this "purely intellectual way" that X was wrong, and uttering the words (silently or aloud) "X is wrong" without any conviction whatsoever? Apparently there would be none. But then it can't really be a kind of knowing.

Take the case of the boss who says that he thinks gender should play no role in hiring decisions, and who believes what he says (namely, that he thinks gender should play no role). This constitutes *prima facie* evidence that the boss believes that gender-based discrimination is wrong. But when we look at his actual hiring practices, we find that he *always* hires male applicants over female ones without any hesitation or misgivings, even in cases in which it is completely clear that the female candidate is better qualified. What are we to say about this case? That the boss believes that gender discrimination is wrong, but always acts contrary to his belief? Surely not. The right verdict is that this man is self-deceived: he thinks he has a moral belief that he in fact lacks.

The simplest argument in favor of Claim 1 is therefore as follows: there must be a difference between knowing that something is wrong and paying empty lip-service

to the claim that it is wrong; likewise, there must be a difference between knowing that something is wrong and mistakenly believing that you know it is wrong; unless Claim 1 is true, there need not be any such difference; therefore, Claim 1 is true.¹³

Claim 2 It is possible for someone who knows all facts *not* to desire that genocide not be perpetrated, *not* to desire that torture not be committed, etc.

This echoes Hume's observation that it is "not contrary to reason to prefer the destruction of the whole world to the scratching of my finger."¹⁴ There is no apparent contradiction in the idea of a being who, knowing all about human suffering of all kinds, never sees anything wrong in one person's inflicting suffering on another. There is no reason to think that Carrie could not look upon all of human life with the clinical indifference of a microbiologist observing bacteria on a petri dish. If there were an individual in whom reason was as highly developed as it is in Carrie, one would hope that it did not display such apathy. But it might.¹⁵

Of course, all the facts is a lot of facts, and a lot more than we can clearly imagine ourselves knowing. Maybe any being that really knew *all* the facts, mental as well as physical, would be somehow compelled to desire the non-occurrence of genocide, torture, etc. If there were real doubt on this score, we would have to suspend judgement on Claim 2—and on M1 as well. But while it is true that we cannot completely put ourselves in Carrie's shoes, we should at least be able to agree that it is possible for someone who knows every morally *relevant* fact about a wrong act not to desire its non-performance. The man who steals my umbrella from the hall knows well enough the morally relevant consequences of his deed, even if he acts with utter disregard for my suffering, and without the faintest hint of shame or remorse. More gravely, a person may murder a competitor out of jealousy, in full knowledge of the relevant consequences of his act, even if he has no desire whatever not to go through with it; indeed, the prospect of these consequences may be exactly what motivates him.

To summarize my objection to M2: that Carrie knows every fact (and thus, by Sect. 1, every truth) is consistent with her failing to desire, for example, that genocide not be perpetrated. But if Carrie does not have this desire, then she does not know that it is wrong to perpetrate genocide. (This is so even if the fact that it is wrong to perpetrate genocide is some natural fact about genocide that Carrie knows.) So, the fact that Carrie knows every fact—including every moral fact, if there are moral facts—is consistent with her failing to know that it is wrong to perpetrate genocide. But if she does not know that it is wrong to perpetrate genocide, then she suffers from moral ignorance. M2 is therefore false: even if all facts are natural facts, someone could know all the natural facts, yet suffer from moral ignorance.

¹³ For a different and much-discussed argument for Claim 1, see Smith (1994, pp. 71–76).

¹⁴ Hume (1978/1739, p. 416).

¹⁵ This is why proponents of ideal observer theories of moral value find it necessary to endow the ideal observer with something more than psychophysical omniscience: see Firth (1952, pp. 333–345), Hare (1981, pp. 44–45), and Smith (1994, pp. 154–161).

3 The moral problem

How should we express the proposition known by someone who has moral knowledge? Suppose we think of a proposition as a set of possible worlds (intuitively, the worlds at which the proposition is true). One might stand in one knowledge relation to this proposition (guised in a certain way) while failing to bear some other knowledge relation to the same proposition (guised in the same way). On the assumption that utilitarianism is true, the proposition known to someone who knows that torture is wrong is the same as the one that is known to someone who knows that it does not maximize utility. We can express this proposition equally with the sentences “Torture is wrong” and “Torture does not maximize utility.” If we want to attribute knowledge of this proposition to someone without saying anything about which knowledge relation he bears to it, we can use an explicitly noncommittal *de re* locution; for example, we can say that the person in question believes, *of* the proposition that torture is wrong, that it is true (or, what comes to the same, that he believes, *of* the proposition that torture fails to maximize utility, that it is true).

In attributing a moral belief to someone, you attribute to him belief of some proposition, but you also attribute to him something more: a desire, as the case may be, or an absence thereof. But that does not mean that there are mysterious propositions with the power to compel anyone who believes them to have (or lack) certain desires. It only means that when we say that someone believes that so-and-so is wrong, what we mean is, among other things, that he desires that so-and-so not be done. We could ascribe belief of the same proposition to the same person without implying his possession of any particular desire (e.g., we could describe him as believing that torture fails to maximize utility, rather than as believing that torture is wrong).

This has a direct bearing on the so-called “moral problem” of metaethics. The moral problem is posed by the following seemingly inconsistent triad¹⁶:

Motivational Internalism: Necessarily, if someone believes that X is wrong, then he desires that X not be done.

Objective Moral Realism: There are moral facts (e.g., of the form *X is wrong*) which are objective, in the sense that they obtain regardless of whether any single individual wants or believes them to obtain.

Queernessless: There is no objective fact that one cannot believe without having a corresponding desire; that is, no objective fact is such that there is a desire such that in order to believe that objective fact, one must have that desire.

That this triad is inconsistent is supposedly established by the following reasoning:

1. If motivational internalism is true, then there are things one cannot believe without having a corresponding desire (For instance, that slavery is wrong is, according to internalism, something one cannot believe unless one desires that slavery not be practiced.).
2. If objective moral realism is true, then what one believes in believing that, e.g., slavery is wrong is an objective fact.

¹⁶ See Smith (1994, pp. 12, 119–125) and Mackie (1977, pp. 38–42).

3. So, if realism and internalism are true, there are objective facts that one cannot believe without having a corresponding desire. (1, 2)
4. But if queernessless is true, then there are no objective facts that one cannot believe without having a corresponding desire.
5. Therefore, queernessless is incompatible with the conjunction of motivational internalism and objective moral realism. (3, 4)

We should reject the first premise of this argument. Motivational internalism entails that there are beliefs that one cannot have without having corresponding desires; for example, it entails that I cannot have the belief that slavery is wrong without desiring that slavery not be practiced. But motivational internalism does *not* entail that there are things that one cannot believe without having corresponding desires; for example, it does not entail that the fact that slavery is wrong is something I cannot believe unless I have a desire that slavery not be practiced.

For the sake of illustration, suppose once more that utilitarianism is true. Then the fact that slavery is wrong is just the fact that slavery does not maximize utility. This, we may suppose, is an objective fact. So, if someone believes that slavery is wrong, what he thereby believes is a certain objective fact. And if someone believes this objective fact, then he *believes the fact* that slavery is wrong. But, given motivational internalism, this does not imply that he *believes that slavery is wrong*. This is because, given internalism, believing a moral fact is only necessary, and not sufficient, for having a moral belief. In addition to believing the fact that slavery is wrong (a.k.a., on our assumption, the fact that slavery fails to maximize utility), one must, in order to have the belief that slavery is wrong, have a certain desire, namely: the desire that slavery not be practiced. And this is a desire that one could lack, even if one knew that slavery failed to maximize utility (or that slavery had whatever other objective property a moral naturalist might want to equate with wrongness).

If motivational internalism is true, the existence of *beliefs that cannot be had* except by those who possess corresponding desires does not entail the existence of *things that cannot be believed* except by those who possess corresponding desires. Internalism therefore does not force its proponents to choose between a subjectivist or antirealist metaethic on one hand, and an ontology of “queer” or “magnetic” states of affairs on the other.

4 The anti-physicalist argument

What about the anti-physicalist argument? Can we mount an old relatum/new relation objection to this as well? Here, the key is to identify necessary conditions of phenomenal knowledge that are not necessary conditions of knowing all facts (or truths). Suppose we can find some condition, P, that, as a matter of logical necessity, anyone satisfies who knows what it is like to have phenomenally red experience, but that someone who knows all the facts could fail to satisfy. Then we may argue that there is no difference between (1) knowing what it is like to have phenomenally red experience, and (2) knowing some purely physical fact while satisfying P. We could then account for Mary’s phenomenal ignorance by observing that she does not satisfy P, without having to suppose that there is any fact she does not know.

Let us say that a person is “experienced in red phenomenology” if and only if he is having red phenomenology, or has had red phenomenology in the past. (Similarly, we can speak of someone as experienced in loud, or sweet, or painful phenomenology.) To a first approximation, being experienced in red phenomenology is a necessary condition of knowing what it is like to have red phenomenology—a necessary condition, that is, of possessing the phenomenal knowledge that **this** (looking at something red) is what it is like to have phenomenally red experience. I say “to a first approximation,” for the following two reasons.

First, suppose that while I am in a deep, dreamless sleep, a perfect, living physical copy of me materializes in the bed next to me. I myself am experienced in red phenomenology, and I know what it is like to have red experience. My sleeping duplicate is not experienced in red phenomenology, but, arguably, he knows as well as I do what it is like to have red experience.

Second, someone might hold that if a person has had a suitably varied range of color experience, and possesses suitable imaginative powers, she may be counted as knowing what it is like to experience red, even if she has never actually had red experience. Such might be the case with a person who has encountered various shades of orange and purple, but not red.¹⁷

These cases are controversial, but I need not take a stand on them one way or another here. Let us call anyone who is experienced in red phenomenology, or resembles someone who is experienced in red phenomenology in suitable structural or functional respects (such as those in which the individuals featuring in the preceding examples resemble ordinary subjects of phenomenally red experience) *virtually experienced in red phenomenology*. I then make the following claim:

Claim 3 Necessarily, if you know what it is like to have red experience, then you are virtually experienced in red phenomenology.

It is hard to see how a proponent of the anti-physicalist argument could object to this claim. The only reason we are given to believe that Mary does not know what it is like to experience red is that she has never had red experience, and does not resemble anyone who has had red experience in the ways that make for virtual experience in red phenomenology. Moreover, if there were, *per impossibile*, a way to know what it is like to have red phenomenology without being virtually experienced in red phenomenology, then, for all we know, Mary must have known in this mysterious way what it was like to have red phenomenology, in which case we would have no reason to accept P1.

My response to the anti-physicalist argument also requires the following claim:

Claim 4 Mary’s knowing every fact is consistent with her not being virtually experienced in red phenomenology.

It is here, presumably, that a proponent of the anti-physicalist argument is apt to dig in his heels; he will hold that, in order to know *every* fact, one must know certain facts that one can know only by having phenomenally red experience (or resembling someone who has had red experience in suitable functional or structural respects).

¹⁷ For a case of the first sort, see Unger (1966, p. 50), and for a case of the latter (Hume 1978/1739, pp. 5–6).

But what is the argument for this? That becoming experienced in red phenomenology normally (and in all real cases) involves learning some previously unknown fact? Suppose that living together with Mary in the grayscale prison is just an ordinary person, call him Larry. When Larry encounters a red object for the first time, he not only becomes virtually experienced in red phenomenology, but also comes to know, of some neural state and of some functional property, that the former has the latter. Larry, who we may suppose spent his days in the prison reading black and white comic books, had no way to specify the relevant neural state or its relevant functional characteristics, prior to actually entering the state and being able to refer to it as the state he is in. But none of this implies that *Mary*, with her extraordinary knowledge of the world, learns a fact previously unknown to her when she has her first experience of red.¹⁸

There is a strong pre-theoretical intuition that there are items of knowledge that one cannot have unless one is virtually experienced in red phenomenology; call this the *virtual experience intuition*. The new-relation objection to P2 respects this intuition. But there are two ways in which the intuition can come out true.

On one hand, the virtual experience intuition might be true because, as the present objection to Claim 4 maintains, there is a certain fact that you cannot know unless you are virtually experienced in red phenomenology.

On the other hand, it might be that the intuition is true because, as I maintain, being virtually experienced in red phenomenology is part and parcel of knowing what red experience is like.

While there is a strong pre-theoretical intuition that at least one of these two claims is true, there is no such intuition that the first in particular is true. What is intuitive is just that since Mary is not virtually experienced in red phenomenology, she doesn't know what it is like to have red experience. It is not intuitive—at least, not pre-theoretically—that since Mary has never had red experience and does not relevantly resemble anyone who has, there is a fact that she does not know.

This, however, leads to a second objection. A dualist may claim that it is pre-theoretically obvious that there is something—some truth, some knowable—that Mary

¹⁸ The contrast between Mary and Larry is instructive in another way. I have argued that having a certain experiential history (or suitably resembling someone who does—a qualification I leave tacit for the remainder of this note) is necessary for possessing phenomenal knowledge, such as knowledge of what it is like to experience red. But it may be asked *why* experience in red phenomenology is necessary for knowing what it is like to experience red. This amounts to asking why we have a verb, “to know,” that is correctly applied only to individuals with appropriate experiential histories. The explanation no doubt lies in the fact that having or lacking a certain experiential history normally—not logically necessarily, but in any real case—goes hand in hand with knowledge or ignorance of a corresponding truth. Larry's lack of colorful experience, coupled with his lack of Laplacean omniscience, puts definite limits on his ability to get things done. When the two are released, Mary will already be able to obey commands such as “Point to a red object” without skipping a beat, since she already knows the microphysical properties of the objects in her environment, the wavelengths of light reflecting off of them and into her eyes, and the names these take in English. This much is a straightforward consequence of her knowing all the physical facts. Larry, by contrast, will be at a total loss, until and unless we teach him the names of the colors. Since the fact that a person has had the conscious experiences he has is a highly reliable (although not, as the possibility of Mary demonstrates, logically infallible) indication that he knows certain facts that he would not otherwise know, it is no surprise that our everyday use of the verb “to know” tracks, among other things, individuals' phenomenal histories.

does not know, given that she lacks virtual experience in red phenomenology. It is at this point that the comparison with the anti-naturalist argument proves especially useful. Is it really obvious that a lack of virtual experience in red phenomenology is possible only on the part of someone who fails to know some fact? No more obvious than that apathy about genocide is possible only on the part of someone who fails to know some fact. The truth is that neither of these claims is obvious—neither is the sort of thing of which we can say, *à la* Moore, that it is more certain than any philosophical argument. This is especially so once it is recognized that rejecting these claims is consistent with agreeing—even insisting—that the absence of a relevant desire or (virtual) experiential history is sufficient for an absence of moral or phenomenal knowledge.

Maybe the best conclusion to draw at this point is not that the second premise of Jackson's argument is false, but merely that, so far, we have no good reason to believe that it is true. If the dualist can provide a reason to think that (virtual) phenomenal inexperience necessitates the existence of a fact unknown to the inexperienced individual, then this will block the present response to the anti-physicalist argument. The prospects for accomplishing this seem dim, short of providing an altogether new argument against physicalism. In any event, it is up to the dualist to make the next move.

To summarize my objection to P2: that Mary knows every fact (and thus every truth) is consistent with her failing to be virtually experienced in red phenomenology. But if Mary is not virtually experienced in red phenomenology, then she does not know what it is like to have red experience. This is just a conceptual truth, which we can also state in terms of how one must relate to a given state of affairs in order to count as knowing what red experience is like: in order for Mary to know what it is like to have red experience, she must bear to some state of affairs—the state-of-affairs that makes it true that **this** is what it is like to have red experience—a knowledge relation that a person bears to a state of affairs only if she is virtually experienced in red phenomenology. Call this relation *R*. The fact that Mary knows every fact is consistent with her failing to know what it is like to have red experience, since it is consistent with her failing to bear *R* to whatever state of affairs it is that makes it true that **this** is what it is like to have red experience. But if she does not know what it is like to have red experience, she suffers from phenomenal ignorance. Therefore P2 is false, or at least doubtful: for all we have reason to believe, even if all facts are physical facts, someone could know all the physical facts, yet suffer from phenomenal ignorance.

5 New guises versus new relations

Both the new-guise and the new-relation strategy exploit the so-called opacity of knowledge. This opacity manifests itself in language as the potential for a true knowledge attribution to be converted into a false one by the substitution of its complement with a different complement having the same truth-conditional content. This happens, for example, when we substitute “there was water beyond the British Isles” with “there was H₂O beyond the British Isles” in the knowledge attribution: “Caesar knew that there was water beyond the British Isles.” Language aside, the opacity of knowledge

consists, at bottom, of the possibility for a person to know that p without knowing that q , even if it is a necessary truth that p iff q .

The difference between the new-guise strategy and the strategy I favor corresponds to a difference over how to account for this possibility. Underlying the new-guise strategy is the idea that whenever it is possible for a single person, at a single time, to know that p without knowing that q , there must be *some* difference between the object of the person's knowledge that p , and what would be the object of his knowledge (if he had it) that q . Underlying the new-relation strategy is the idea that a person might know that p without knowing that q even if the object of the knowledge that p is one and the same as the object of the knowledge that q . On this view, the difference between knowing that p and knowing that q may just be the difference between bearing one versus another knowledge relation to a *single* fact, proposition, or truth (the truth that p , a.k.a. the truth that q).¹⁹

It is important to see that the new-relation strategy is immune to the objection raised against the new-guise strategy in Sect. 1. On the new-relation view, the fact that Mary knows every knowable is entirely consistent with her failing to know what it is like to have red experience, since it is entirely consistent with her never having had red experience (or suitably resembling someone who has had such experience). It is a fact that Mary *would* know what it was like to have red experience *if* she were virtually experienced in red phenomenology, and this is a fact that Mary knows even prior to her release. But from this it does not follow that she knows what it is like to have red experience. Unlike the new-guise strategist, a proponent of the new-relation approach is not compelled to place any fact out of the imprisoned Mary's cognitive reach, in order to account for the enlargement that her truth-apt knowledge undergoes upon her release.

To reinforce the point that there is a genuine difference between the new-guise and new-relation strategies, it helps to consider some approaches to the problems we have been considering that, while clearly intended to employ *one* of the strategies, leave it unclear which. The uncertainty that surrounds these approaches shows that there really are two strategies to choose from here.

5.1 Loar's approach to the anti-physicalist argument

First, consider Brian Loar's response to the anti-physicalist argument. This agrees with my response in several important respects. We agree that (1) phenomenal knowledge is a kind of truth-apt knowledge or "knowledge-that" (as opposed to "knowledge-how," as Nemirow, Lewis, and others have suggested). We also agree that (2) when Mary leaves her prison, she acquires new truth-apt knowledge. Finally, we agree that (3) it does not follow from (2) that the imprisoned Mary's ignorance of what it was like to experience red involved the existence of a fact or state of affairs unknown to her while imprisoned.²⁰

¹⁹ I have developed a generalized old relatum/new relation account of epistemic opacity elsewhere; see Pelczar (2007).

²⁰ See Loar (1990).

The uncertainty I wish to point out concerns Loar's reasons for holding (3). For a person to know something is for a certain three-part state of affairs to obtain: the first part is the knower, the last is what is known (call this "the knowable"), and in between these is the relation that the knower bears to the knowable so as to count as having the knowledge in question. This means that there are two ways in which Mary can acquire new knowledge. One is by coming to bear a relation of knowledge to a new knowable—a knowable to which she did not previously bear any relation of knowledge; call this "the First Way." The other is by coming to bear to an old knowable—a knowable to which she already bore some knowledge relation—a knowledge relation that she did not previously bear to that knowable; call this "the Second Way."

I maintain that (3) is true because Mary learns upon her release only by acquiring knowledge in the Second Way. In particular, she learns inasmuch as she comes to bear to some knowable (which might as well just be some physical-functional state of affairs) a relation of knowledge that has among its instantiation conditions that the one who bears it to a given state of affairs have had reddish phenomenology (or relevantly resemble someone who has had such phenomenology).

Now, what about Loar? Does he think that Mary acquires knowledge in the First Way, or only in the Second Way? If in the First Way, he is a proponent of the old fact/new guise strategy; if only in the Second Way, he is a proponent of the old relatum/new relation approach.

Unfortunately, it is difficult to tell which way Loar intends to go. Some of his comments suggest that he endorses the old fact/new guise strategy. For example, he writes:

Knowing how a state feels is knowing that it feels a certain way. Anti-physicalists are right in holding that it is the possession of distinctive information, for it involves a genuinely predicative component of judgment—what I am calling a phenomenal concept—whose association with any physical-functional concept is straightforwardly a posteriori.²¹

On a natural reading, to say that "knowing that a state feels a certain way is to possess distinctive information" is to say that it is to know some distinctive knowable. But information-talk is somewhat fluid, and it is possible that Loar would want to say that, just as there are two ways of gaining knowledge, there are also two ways of acquiring information (corresponding to Ways One and Two, described above). Again, immediately after the passage just quoted, he writes, with reference to the ability objection to the anti-physicalist argument:

A physicalist would be forced into the Nemirow/Lewis reply if he were to individuate bits of knowledge, cognitive information, in terms of possible-world truth conditions.

Here again I agree with Loar on one interpretation, but not on another. I maintain that it is possible to individuate *knowables* as sets of possible worlds (or even more coarsely

²¹ Loar (1990, p. 85).

than that) without being forced into the Nemirow/Lewis reply (or any other). So, if by “bits of knowledge” and “cognitive information” Loar means knowables, I disagree with his remark. But if by “bits of knowledge” and “cognitive information” he means states, episodes, or cases of knowing, then I agree with him. To get a truly satisfying response to the anti-physicalist argument, we must individuate an individual’s knowledge states not just on the basis of their objects (whether these be sets of possible worlds, guised propositions, or anything else) but also on the basis of the knowledge relations the states involve.

At the end of the day, Loar’s comments leave it unclear whether he favors a First Way (old fact/new guise) or Second Way (old relatum/new relation) response to the knowledge argument. This in turn reinforces the point that there is something to choose between these two responses, and that they should not be conflated.²²

5.2 Sturgeon’s approach to the moral problem

In a recent essay, Scott Sturgeon proposes a solution to the moral problem that in some ways resembles my own.²³ Sturgeon and I agree that the key to resolving the apparent conflict among Motivational Internalism, Objective Moral Realism, and Queernessless lies in the correct individuation of beliefs. In particular, we agree that different belief states (of the same agent at the same time even) can be alike in their truth-conditions. A classification of an agent’s belief states according to their truth-conditions may therefore yield a coarser individuation than a classification that takes a difference of truth-conditions as perhaps sufficient, but not necessary, for a difference in belief states.

So much is common ground between Sturgeon and me. If we disagree, it is over the exact way in which a finer individuation of belief states stands to resolve the moral problem.

For Sturgeon, the finer individuation has to do with what he calls “normative concepts”: two beliefs with the same truth-conditions might differ due to the use

²² Loar is not alone in giving an account that wavers between the First and Second Ways. In Bigelow and Pargetter (1990, pp. 138–144), John Bigelow and Robert Pargetter point to the possibility of a Second Way response to the anti-physicalist argument; but then they seem to revert to the First Way when they say that Mary’s first color experience “enables her to discriminate among new representable possibilities.” (Bigelow and Pargetter 1990, p. 144) Similarly, when John Perry says that upon having her first colorful experiences, Mary’s “beliefs will change...[b]ut the demands that the truth of her beliefs place on the world do not” (Perry 2001, p. 112), he seems to be proposing a Second Way response to Jackson’s argument. But then he seems to revert back to the First Way, when he writes that “[i]n Mary’s case...the need is not for nonphysical properties, but for a broader conception of the content of thought,” where the “content of a belief is simply whatever is believed about whatever the belief is about.” (Perry 2001, p. 113) Terence Horgan says that Mary learns upon having her first experience only inasmuch as she becomes acquainted with phenomenal redness from a new “experiential” perspective (Horgan 1984, p. 151); this is suggestive of the Second Way. But in the same place, Horgan says that acquiring this new perspective results in Mary’s adding to her stock of ontologically physical information, which is suggestive of the First Way—a suggestion that is reinforced by Horgan’s emphasis on an analogy between the imprisoned Mary’s failure to know what it is like to experience red, and Lois Lane’s failure to know that Clark Kent can fly (Horgan 1984, pp. 150–152).

²³ Sturgeon (2007).

or appearance in one, but not the other, of a normative concept. Consider someone, call him Paul, who believes whatever fact it is that makes it true that genocide is wrong, but who does not have the moral belief (and does not make the normative judgement) that genocide is wrong. (On Sturgeon’s view, as on mine, a person like Paul is perfectly possible.) Now suppose that Paul undergoes a change of outlook, so that he does come to believe that genocide is wrong. By Sturgeon’s account, this change amounts to Paul’s coming to believe anew some fact that he already believed, by deploying a normative concept that he did not previously deploy in this way. To determine whether Sturgeon advocates a new-guise or a new-relation approach, we must figure out exactly wherein lies the difference between two beliefs that differ only in that one, but not the other, involves the use of a normative concept.

As Sturgeon describes it, the use of a normative concept essentially involves the presence of a desire; as he puts it, a normative concept gets “realised” in a judgement only in the company of a desire concerning the non-normative content of the judgement.²⁴ The appeal to “non-normative content” suggests that Sturgeon thinks that the presence of a normative concept has a bearing on the content of belief or judgement, rather than on the believing or judging relation. On this view, two beliefs that differ only in that one but not the other manifests a normative concept are beliefs that differ in their *objects*: the object of the belief that manifests the normative concept is the result of attaching that normative concept to the non-normative content of the belief that does not manifest the concept. This would be a version of the new-guise theory of normative belief.

On the other hand, it is hard to interpret Sturgeon’s claim that

the concept *OUGHT* used in your judgement requires for its manifestation (in such judgement) a *bona fide* desire concerning its non-normative content²⁵

as anything more than a roundabout version of my claim that

in order to count as believing that so-and-so is wrong, you must have a desire that so-and-so not be done.²⁶

If this is what Sturgeon’s position comes to, it seems he should simply adopt the old relatum/new relation strategy which I favor. For in this case, it is the presence or absence of the relevant desire that does all the work; “normative content” drops out of the picture entirely.

The important point for my purposes is that it is unclear whether Sturgeon’s proposed solution to the moral problem goes by way of an old fact/new guise strategy, or an old relatum/new relation strategy. Once again, the uncertainty demonstrates the need to keep the two strategies distinct. This is all the more so in view of the special difficulties that face the old fact/new guise approach.

²⁴ Sturgeon (2007, pp. 569, 584).

²⁵ *Ibid.*, p. 580.

²⁶ Or, lack a certain desire, if the judgement is to the effect that such-and-such is right.

6 Conclusion

The imprisoned Mary does not know what it is like to have red experience; that is, she does not satisfy both of the following conditions: (1) that of being virtually experienced in red phenomenology, and (2) that of knowing a relevant (maybe physical) fact of visual psychology. As it happens, Mary does satisfy condition (2); her ignorance of what it is like to experience red results from her failure to satisfy (1). In this she resembles Carrie, whose moral ignorance does not arise from there being some fact unknown to her, but only from her failing to have certain relevant desires, such as the desire that acts of torture not be committed. Mary's phenomenal ignorance concerns a neuro-functional fact, and results from her neither having had phenomenally red experience, nor relevantly resembling anyone who has. Carrie's moral ignorance concerns facts that it is the task of normative ethics to describe, and results from her lacking certain desires. In neither case does the compatibility of the posited ignorance with knowledge of all the physical or natural facts cast doubt on the idea that the world is a completely natural, and indeed completely physical, place.

References

- Bigelow, J., & Pargetter, R. (1990). Acquaintance with qualia. *Theoria*, 56, 129–147.
- Boyd, R. (1988). How to be a moral realist. In G. Sayre-McCord (Ed.), *Essays on moral realism* (pp. 181–228). Ithaca: Cornell University Press.
- Brink, D. O. (1989). *Moral realism and the foundations of ethics*. Cambridge: Cambridge University Press.
- Brink, D. O. (2001). Realism, naturalism, and moral semantics. *Social Philosophy and Policy*, 18(2), 154–176.
- Broad, C. D. (1925). *The mind and its place in nature*. London: Kegan Paul, Trench, Trubner & Co.
- Chalmers, D. J. (2004). Phenomenal concepts and the knowledge argument. In P. Ludlow, Y. Nagasawa, & D. Stoljar (Eds.), *There's something about Mary: Essays on Frank Jackson's knowledge argument against physicalism* (pp. 269–298). Cambridge: MIT Press.
- Dreier, J. (1992). The supervenience argument against moral realism. *Southern Journal of Philosophy*, 30(3), 13–38.
- Firth, R. (1952). Ethical Absolutism and the ideal observer. *Philosophy and Phenomenological Research*, 12(3), 317–345.
- Frankena, W. K. (1939). The naturalistic fallacy. *Mind*, 48(192), 464–477.
- Hare, R. M. (1981). *Moral thinking: Its levels, method, and point*. Oxford: Clarendon Press.
- Harman, G. (1977). *The nature of morality: An introduction to ethics*. New York: Oxford University Press.
- Horgan, T. (1984). Jackson on physical information and qualia. *The Philosophical Quarterly*, 34(135), 147–152.
- Hume, D. (1978/1739). *A treatise of human nature*. Oxford: Clarendon Press.
- Jackson, F. (1982). Epiphenomenal qualia. *The Philosophical Quarterly*, 32(127), 23–36.
- Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy*, 83(5), 291–295.
- Jackson, F. (2005). Consciousness. In F. Jackson & M. Smith (Eds.), *Oxford handbook of contemporary philosophy* (pp. 310–333). Oxford: Oxford University Press.
- Lewis, D. (1990). What experience teaches. In W. Lycan (Ed.), *Mind and cognition* (pp. 499–518). Oxford: Blackwell.
- Loar, B. (1990). Phenomenal states. *Philosophical Perspectives IV: Action Theory and Philosophy of Mind*, 81–108.
- Lockwood, M. (1989). *Mind, brain and the quantum: The compound 'I'*. Oxford: Blackwell.
- Lycan, W. G. (1986). Moral facts and moral knowledge. *Southern Journal of Philosophy*, 24(Supplement), 79–94.
- Lycan, W. G. (1990). What is the “subjectivity” of the mental? In J. Tomberlin (Ed.), *Philosophical perspectives IV: Action theory and philosophy of mind* (pp. 109–130). Ridgeview: Atascadero.

- Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. New York: Penguin Books.
- Miller, R. W. (1985). Ways of moral learning. *Philosophical Review*, 94(4), 507–556.
- Moore, G. E. (1903). *Principia ethica*. New York: Cambridge University Press.
- Nemirow, L. (1980). Review of mortal questions, by Thomas Nagel. *The Philosophical Review*, 89(3), 473–477.
- Pelczar, M. (2005). Enlightening the fully informed. *Philosophical Studies*, 126(1), 29–56.
- Pelczar, M. (2007). Forms and objects of thought. *Linguistics and Philosophy*, 30(1), 97–122.
- Perry, J. (2001). *Knowledge, possibility, and consciousness*. Cambridge: MIT Press.
- Putnam, H. (1975). Language and reality. *Mind, language, and reality* (pp. 272–290). Cambridge: Cambridge University Press.
- Quine, W. V. O. (1960). *Word and object*. Cambridge: MIT Press.
- Sayre-McCord, G. (1997). ‘Good’ on twin earth. *Philosophical Issues: Truth*, 8, 267–292.
- Smith, M. (1994). *The moral problem*. Oxford: Blackwell.
- Sturgeon, N. (1988). Moral explanations. In G. Sayre-McCord (Ed.), *Essays on moral realism* (pp. 229–255). Ithaca: Cornell University Press.
- Sturgeon, S. (2007). Normative judgement. *Philosophical Perspectives*, 21, *Philosophy of Mind*, 569–587.
- Thau, M. (2002). *Consciousness and cognition*. Oxford: Oxford University Press.
- Unger, P. (1966). On experience and the development of the understanding. *American Philosophical Quarterly*, 3(1), 48–56.
- Wittgenstein, L. (1965/1930). A lecture on ethics. *The Philosophical Review*, 74(1), 3–12.