

Consequentialism, Metaphysical Realism, and the Argument from Cluelessness[†]

Dale Dorsey

Department of Philosophy
University of Kansas
1445 Jayhawk Boulevard
Wescoe Hall, rm. 3090
Lawrence, KS 66045
ddorsey@ku.edu

ABSTRACT: The most powerful version of the classic epistemic argument against consequentialism is stated in an article by James Lenman. Lenman’s “argument from cluelessness” claims that a significant percentage of the consequences of our actions are wholly unknowable and hence, when it comes to assessing the moral quality of our actions, we are literally without a clue. In this paper, I distinguish the argument from cluelessness from traditional epistemic objections to consequentialism, and argue that the argument from cluelessness should be no more problematic for consequentialism than the argument from epistemological skepticism should be for metaphysical realism. I argue that this puts those who would reject consequentialism on grounds of cluelessness in an awkward philosophical position.

The most powerful version of the classic epistemic argument against consequentialism is stated in an article by James Lenman.¹ Lenman’s argument—henceforth known as the “argument from cluelessness”—claims that a significant percentage of the consequences of our actions are wholly unknowable and hence, when it comes to assessing the moral quality of our actions, we are literally without a clue. If we accept consequentialism, the moral valence of our actions is simply unknowable; if that is correct, it is difficult to see how the moral valence of actions and decisions could possibly play any meaningful role in human life.

This argument is powerful. Indeed, I shan’t dispute the claim that many of the consequences of our actions are unknowable to us. The question I shall address is whether the fact of cluelessness should worry the advocate

[†]Thanks to Aaron Garrett for inspiring my attention to this important issue. I would also like to dedicate this paper to the memory of my late colleague, A.C. (Tony) Genova, whose early encouragement led to its first complete draft.

¹“Consequentialism and Cluelessness” in *Philosophy and Public Affairs* 29 (2000).

of consequentialism. I claim it should not. In particular, I argue that there is very good reason to believe that the problem of cluelessness for consequentialism is only as embarrassing for consequentialism as the spectre of epistemological skepticism is embarrassing for metaphysical realism. But coping with epistemological skepticism rarely tempts anyone to abandon metaphysical realism. Moral philosophers should face no greater temptation when it comes to consequentialism. Or, at any rate, so I shall argue.

1. *Epistemic Arguments and the Argument from Cluelessness*

Epistemic arguments against consequentialism generally run like this: because consequentialism determines the moral valence of any action by that action's consequences, we have epistemic access to the moral valence of any action if and only if we have epistemic access to the consequences of said action. But because we do not know, *ex ante*, the consequences of any particular action, we cannot know that action's moral valence. But if we cannot know the moral valence of our actions in the *ex ante*, then we cannot know which actions consequentialism requires us to perform. But for consequentialism to offer meaningful guidance to moral agents, moral agents must know, in the *ex ante*, which actions consequentialism requires them to perform. Hence consequentialism cannot offer meaningful guidance to moral agents. Consequentialism cannot adequately conduct the "very business of ethics".²

One common reaction to the epistemic argument is to admit fallibility in the *ex ante*, but to insist that consequentialism can offer meaningful guidance given the past results of actions of similar types. Indeed, this was Mill's response:

During [the entire duration of their species], mankind have been learning by experience the tendencies of actions; on which experience all the prudence, as well as all the morality of life, are dependent. . . It is truly a whimsical supposition that, if mankind were agreed in considering utility to be the test of morality, they would remain without any agreement as to what is useful, and would take no measures for having their notions on the subject taught to the young, and enforced by law and opinion.³

One might put this point in the following way. Insofar as I'm uncertain about

²Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection" in *Ethics* 101 (1991), 466-7.

³John Stuart Mill, *Utilitarianism* Chapter II, para. 28-29.

the consequences of any particular action, consequentialism cannot offer any direct *ex ante* guidance of action. Nevertheless, this does not mean action-guidance is barred to the consequentialist. Consequentialism can admit *ex post* action-guidance. For Mill, we *learn*, given our own experience and the experience of others, what sort of actions have what sort of consequences. Given this knowledge, we might formulate specific decision procedures on the basis of backward-looking assessments of actions. For instance, one might learn, given this backward-looking assessment, that murders generally turn out to have bad consequences. Insofar as murders generally have bad consequences, we will attempt to avoid murder when making moral decisions. But this decision is based entirely on the consequentialist assessment of the moral valence of past instances of murder. Hence consequentialism can allow for meaningful, though *ex post*, action-guidance.

Of course, this still leaves a certain degree of *ex ante* uncertainty with regard to the consequences of any particular action. The future may be unlike the past. Perhaps the ways in which possible actions are relevantly different from similar past actions may be inaccessible or indeterminate given our decision procedures. Necessarily, because the decision procedure is designed with only knowledge of the moral valence of past actions, that it will be successful in guiding moral agents to the best action is, at best, probabilistic. Nevertheless—and this is the second part of the standard response to the epistemic argument—this form of uncertainty should not be feared. To insist that the moral valence of all actions must in every case be knowable in the *ex ante* is to insist on a standard too high to represent our common-sense moral experience. To insist, in other words, that the epistemically responsible agent must have full knowledge of the moral valence of all courses of action open to her is to insist on a mistaken standard of moral epistemic access. But because consequentialism allows *ex post* access to the consequences of actions, we have at least a reasonable body of evidence to use in making *ex ante* decisions. Sometimes, however, we come to learn that what we thought was the right thing to do was, in fact, the wrong thing despite our best efforts. Such is life.

For the purposes of argument, I will assume that this response to the traditional epistemic argument succeeds (as Lenman himself seems to⁴). However, the argument from cluelessness is not an epistemic argument of the traditional form, and is not susceptible to the above response. The method of responding to *ex ante* uncertainty by means of *ex post* knowledge of consequences obviously requires *ex post* access to the consequences of in-

⁴See Lenman, 360.

dividual actions.⁵ But Lenman's claim is not simply that consequentialism can offer no *ex ante* guidance to deliberating moral agents (though his argument certainly implies this also). Rather, Lenman claims that *no one* can have *any epistemic access* to the moral valence of *any action whatsoever*, ever. If consequentialism is true, for any action ϕ , the moral valence of ϕ (whether ϕ is past, present, or future) is totally unknowable, beyond the epistemic access of anyone.

To introduce a bit of terminology, one might divide the consequences of any given act into three different potential epistemic categories:

1. Visible, certain consequences: certain consequences of a given action ϕ are those consequences of ϕ that can be foreseen, in the *ex ante*, by an epistemically responsible agent.
2. Visible, uncertain consequences: uncertain consequences of a given action ϕ are actions that could not have been foreseen, in the *ex ante*, by an epistemically responsible agent, but could be ascertained in the *ex post*.
3. Invisible consequences: invisible consequences of a given action ϕ are actions that could not have been foreseen in the *ex ante*, nor ascertained in the *ex post*, by an epistemically responsible agent.

For the traditional response to succeed, the consequences of any particular action must fit into the first two epistemic categories. For consequentialism to offer any action-guidance whatever, there can be no consequences that fall into category (3). But according to Lenman, the problem is not simply that the consequences of actions are visible, but uncertain. Rather, the problem is that the consequences of actions—or, at least, a big chunk of them—are *invisible*, unknowable in the *ex ante* or *ex post*.⁶ For Lenman, *most* of the consequences of our actions fit into category (3).

Why believe this? Two reasons. First, virtually anything we do, as Lenman notes, is likely to have consequences that are *identity-affecting*, altering facts about who exists or will exist.⁷ But the consequences of identity-affecting actions, or most of them anyway, are invisible. As an illustration, Lenman considers the act of “Richard,” an early German conquerer, in sparing “Angie”—whose *far* distant ancestor turned out to be Adolf Hitler:

The decision to spare Angie is an event with massive causal ramifications. It is highly plausible that almost all killings and engen-

⁵See Lenman, 360-1.

⁶See, for instance, Lenman, 360-363.

⁷See Lenman, 347.

derings and refrainings from these have similarly massive causal ramifications. These actions ramify in massive ways most obviously because they are, let us say, ‘identity-affecting’. These are actions that make a difference to the identities of future persons and these differences are apt to amplify exponentially down the generations. A very high proportion of identity-affecting actions are, it is enormously plausible, reliably subject to such massive causal ramification.⁸

Given that any particular action—especially actions that are in some way or other identity-affecting—will have massive causal ramifications, we surely have no good evidence in the *ex ante* or *ex post* about the axiological quality of the consequences of our actions. That Richard’s sparing of Angie eventually resulted in the holocaust is certainly unknowable to Richard, Angie, or, indeed, *anyone else*.

Second, some of the consequences of our actions are the result of that action playing a role in a causal system that is extremely sensitive to small changes. Lenman writes: “Indeed, it is arguably a very real possibility that very many actions that seem very insignificant are subject to massive causal ramification. For some causal systems are known to be extremely sensitive to very small and localized variations or changes in their initial conditions.”⁹ Lenman cites the weather and financial markets, systems that can alter significantly with very small perturbations. These consequences are barred to even the most epistemically responsible agent. Hence a gigantic proportion of the consequences of our actions are invisible, and hence the moral valence of our actions is invisible as well. Furthermore, it would appear that this argument could go further. It needn’t be the case that any particular action *actually* had massive, unknowable causal ramifications for the moral valence of that action to be unknowable. Because we know that any action *can* have such causal ramification, we are unjustified in maintaining any confidence that the visible consequences of any particular action are the only consequences (even if, in fact, they are). Hence the argument from cluelessness seems to show, quite powerfully, that the moral valence of any action is simply unknowable, beyond our epistemic access. As Lenman writes: “The worry is not that our certainty is imperfect, but that we do not have a clue about the overall consequences of many of our actions. Or rather—for let us be precise—a clue is precisely what we do have, but it is a clue of bewildering insignificance bordering on uselessness—like a detective’s discovery of a

⁸Lenman, 346.

⁹See Lenman, 347.

fragment of evidence pointing inconclusively to the murderer's having been seven feet tall. We may not be *strictly* without a clue, but we are *virtually* without a clue.”¹⁰

2. *The Balancing-Out Hypothesis*

Even if we accept that consequentialism can acceptably accommodate the existence of visibly uncertain consequences, the argument from cluelessness shows that epistemic access to the moral valence of our actions is much more limited than the traditional response to epistemic arguments can allow. Because some (perhaps most) consequences of a given action are invisible, consequentialism has no grounds on which to play any meaningful role in human life, whether as a guide to action, or as a method by which to assess past actions.¹¹ But a response is open. The consequentialist is rightly unbothered by the mere *existence* of invisible consequences. As we all know, consequentialism does not determine the moral valence of ϕ on the basis of the mere *number* of ϕ 's consequences. The consequentialist determines the moral valence of ϕ on the basis of the *comparative quality* of ϕ 's consequences—compared, that is, to other available actions. If this is correct, consequentialism implies a lack of epistemic access to the moral valence of actions only if invisible consequences are enough to affect the moral quality of one action compared to another. But if we assume that the comparative quality of invisible consequences (for any two acts ϕ and ψ) will *balance out* in the long run, the mere existence of invisible consequences makes no difference to the relative moral quality of ϕ and ψ . We can know the comparative moral quality of ϕ and ψ by considering only their visible consequences; their invisible consequences will be of (roughly) identical quality. Consider:

The Balancing-Out Hypothesis: For any two actions ϕ and ψ , the difference in axiological quality of the total set of invisible consequences of ϕ and ψ will not be large enough to alter the comparative moral valence of ϕ and ψ as assessed by their visible consequences.

The balancing-out hypothesis is *a priori* plausible, for two reasons. First, given that the causal ramifications of our actions are massive, it seems plausible to say that any bad invisible consequence of ϕ will be matched by an

¹⁰Lenman, 349-50.

¹¹Compare Lenman's discussion of "disengaged consequentialism", Lenman 361-2.

equally bad invisible consequence, perhaps hundreds of years down the line, by any other action ψ . The key point here is the *sheer number* of invisible consequences. Given facts of “massive causal reverberation”, it seems doubtful to believe that the invisible consequences of ϕ (bad or good) won’t be mirrored by equally bad (or nearly equally bad) invisible consequences of ψ .¹² In addition, it seems plausible that the invisible consequences of any individual action ϕ will be unlikely to trend as significantly positive or negative. In other words, if there is a very bad invisible consequence of ϕ , it seems plausible to believe that there is an equally *good* invisible consequence of ϕ ; if so, the total axiological value of invisible consequences of ϕ will trend toward zero. For these reasons, in considering my potential donation of \$100 to Oxfam International, it seems right, at least in the *a priori*, to believe that the sum of the invisible consequences of donating versus not donating are not likely to be enough to outweigh the visible superiority of donating rather than not donating.

The balancing-out hypothesis has been marshalled in response to the argument from cluelessness by Tyler Cowen. Cowen argues that we can have confidence that the invisible consequences of any given set of actions will not be able to outweigh their visible consequences, especially when difference in quality of visible consequences of two actions ϕ and ψ are quite large. Cowen writes:

[C]onsider the value of stopping a terrorist attack that would decimate the entire United States. Or consider an attack that would devastate all of Western civilization, or the entire world. At some point we can find a set of consequences so significant that we would be spurred to action, again in open recognition of broader long-run uncertainties.

Surely at some point the upfront change must be large enough to provide a persuasive reason for or against it. What if a cosmological disaster destroyed 99.9999 percent of all intelligent life across the universe? Yes, it is possible that subsequent cosmological events could lead to an even greater blossoming of wonders, but at some point of comparison this point is simply fatuous. Most of the life in the universe is being destroyed and more likely than not this is a horrible catastrophe even in the much longer run. So we can argue ‘how large’ an upfront event is needed to

¹²Compare the classic response to the epistemic objection offered by J.J.C. Smart, “Outline of a System of Utilitarian Ethics” in *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1974), 33.

sway us toward an evaluative judgment, but a sufficiently large upfront event should do the trick.¹³

Consider the cartoonishly evil action of detonating a nuclear explosion at the center of the earth that will destroy the planet and all life with it, and the action of refraining from so doing. The difference in quality of the visible consequences of these two actions is of such a magnitude that it seems “fatuous” to say that invisible consequences might outweigh it.

The most serious problem with Cowen’s response, and indeed with the balancing-out hypothesis itself, is that even if we confine the balancing-out hypothesis to “big” actions, there remains a serious problem: though the balancing-out hypothesis *seems plausible*, there is *no evidence in its favor*, and hence *no good reason to believe it*, whether or not the actions in question are “small” or “big”. Lenman writes:

The fundamental problem is that we cannot just carve up the various possibilities and assign equal probability to each of them without some basis on which to favour some one of the many ways of carving them up over others. We need some nonarbitrary way to partition the possibilities before us and there very plausibly isn’t one—or rather there isn’t one unless we appeal to the sort of prior information about probabilities which, *ex hypothesi*, we do not have.¹⁴

Because there is no evidence that the balancing-out hypothesis is more likely than the falsity of the balancing-out hypothesis, we are surely not licensed to simply assume that invisible consequences of ϕ -ing and ψ -ing will be of equivalent value. Cowen has precisely *no* data upon which to rest his argument that those actions we take that have “big” visible consequences will balance out their respective invisible consequences. We have no earthly idea *how* big the visible consequences must be to outweigh invisible consequences. The invisible consequences are, after all, *invisible*. Taking again the case of Richard and Angie, Lenman writes: “So once again the chance that the difference between gains and losses is so slight that the intrinsic evil of Angie’s death is enough to bridge it is vanishingly small. That chance is not zero and the difference made by Angie’s death can only be negative, so her death generates a reason against killing her. But it is an *extremely weak*

¹³Tyler Cowen, “The Epistemic Problem Does Not Refute Consequentialism” in *Utilitas* 18 (2006), 387.

¹⁴Lenman, 354.

reason.”¹⁵ According to Lenman, the chance that the invisible consequences of two actions will balance-out sufficient to render their visible consequences determinate of their comparable moral valence is miniscule, and hence we have no reason to accept the balancing-out hypothesis, and hence no reason to believe consequentialism can offer any moral guidance in the *ex ante* or *ex post*.

3. *Invisible Consequences and the Skeptical Scenario*

My response to the argument from cluelessness begins by exploiting a minor error in Lenman’s response to the balancing-out hypothesis. Lenman claims that the “chance that the difference between gains and losses is so slight that the intrinsic evil of Angie’s death is enough to bridge it is vanishingly small.” In another context, Lenman describes the chance that the visible consequences will be enough to tip the balance in favor or against another particular action as “*astronomically* improbable.”¹⁶ But just as it is worth wondering what evidence Cowen, or anyone else, might have *for* the balancing-out hypothesis, it is worth wondering what evidence Lenman has *against* the balancing-out hypothesis. He appears to be making a probability calculation, viz., the probability that any visible consequences will be enough to decide the overall value of all total consequences (including invisible consequences) is “astronomically” improbable. But Lenman has no grounds for making this claim. In fact, the suggestion that the invisible consequences will balance-out, and the suggestion that they will *not* balance-out, are on precisely the same epistemic footing. We have *no idea*, given the fact that invisible consequences are, literally, *invisible*, whether they will or will not be slight enough to allow the visible consequences to determine the moral valence of ϕ in comparison to ψ . Given that the invisible consequences of ϕ are literally invisible, the probabilities are simply indeterminate. To put this point in another way, to make any claim about the balancing-out hypothesis, either for or against, one must have at least some epistemic access to invisible consequences. But the argument from cluelessness positively relies on the claim that we have no such access; without this claim, the argument from cluelessness is just another statement of the more traditional epistemic argument against consequentialism to which consequentialists have a practiced (and, so we’re assuming here, successful) response. Hence, according to the argument from cluelessness itself, we

¹⁵Lenman, 356.

¹⁶Lenman, 356.

can make no claim whatsoever about the likelihood or unlikelihood of the balancing-out hypothesis.

Lenman might reply: “so what?”. It remains the case that we have no epistemic access to the invisible consequences of our actions, and hence have no idea whether any particular action in the *ex ante* or *ex post* actually leads to beneficial consequences overall. Of course, *if* the invisible consequences of alternative actions balance out, consequentialism can allow epistemic access to the moral valence of actions in, at least, the *ex post*. But if these consequences do not balance out, we are simply in the dark about the moral valence of actions. And given that the facts about whether invisible consequences will balance out are *invisible*, we have no way to assess the probability that, in fact, they will balance-out. Hence we should have no confidence that consequentialism allows epistemic access to facts about the moral valence of actions. At this point, it would seem that there is no acceptable consequentialist response to the argument from cluelessness. Whether in the *ex post* or *ex ante*, we are in the dark about the moral valence of actions on a consequentialist view.

One might put the predicament facing consequentialism in the following way. Consequentialism faces, as it were, a *skeptical scenario*: that the balancing-out hypothesis is false. If the balancing-out hypothesis is true, consequentialism has no worries about epistemic access, at least in the *ex post*, to the relative consequences of ϕ versus ψ . But because we cannot know whether the balancing-out hypothesis is true or false (given the fact that these consequences are wholly invisible), we cannot know whether the invisible consequences of our actions balance out, and hence cannot know whether the visible consequences of our actions are sufficient to determine the moral valence of our actions. In short, though the balancing-out hypothesis may *seem* plausible in the *a priori*, we can have absolutely no confidence that it is actually true. And in that case, consequentialism succumbs to a form of epistemological skepticism: *given* that we cannot rule out this skeptical scenario (i.e., that the balancing-out hypothesis is false), the moral valence of our actions is unknowable.

4. *Metaphysical Realism and Epistemological Skepticism*

At this point in the paper, I want to shift gears. Leave aside consequentialism and consider the ontological thesis known as “metaphysical realism”. I will not attempt to give a very precise characterization of metaphysical realism here. For the purposes of this paper, however, by “metaphysical realism” I simply mean the view that it is possible that the best understanding

open to an epistemically responsible agent of the way the world is is not the way the world *actually* is. Metaphysical realism has a number of relevant contrast views, but these contrast views share the claim that facts about the external world are, in some sense or other, determined by what we believe, what our best theories say or will say, or what is in-principle verifiable by reasonable epistemic agents. Realism, on the other hand, drives a modal wedge between our conceptualization of the world, and the world itself.

However, this modal wedge can be exploited by skeptics. Skeptics claim that the possibility of error is sufficient to render knowledge of the external world impossible. Consider, for instance, the infamous “evil demon” hypothesis familiar from Descartes. It is certainly metaphysically possible that an evil demon is deceiving us, in other words, allowing our perceptions of the way the world is to differ substantially from the way the world, in fact, is.¹⁷ A classic variant is the “brain in a vat” scenario.¹⁸ Given our best understanding of the world, we have no ability to rule out the possibility that we might, in fact, be brains in vats, or living in an experience machine, in which our perceptions of the external world are mere simulations. Under such conditions, the world as it “really is” is substantially different than the best understanding of the world open to an epistemically responsible agent.¹⁹ The possibility of this scenario is a direct result of metaphysical realism’s insistence on a distinction between the way the world actually is, and our perceptions or theories of the world.

Whether any or all of these skeptical scenarios holds *is invisible*. We have no idea whether we are brains in vats, or whether we are in a computer simulation, or whether there is an evil demon distorting our perceptions, etc. It might seem as though these skeptical scenarios are not very likely. Hence, one might be tempted to conclude, they are relatively safe to ignore when it comes to gaining knowledge about the world around us. But this inference is unlicensed. Because whether a skeptical scenario holds is simply invisible, we have no basis whatever to infer its probability. Though it seems to me very likely that I am not being deceived by an evil demon, I simply have no legitimate basis for making this claim—this fact is, in essence, the very heart of the skeptical challenge. Given that every skeptical scenario is by its very nature invisible, we lack sufficient epistemic access to

¹⁷See Rene Descartes, *Meditations on First Philosophy*, first meditation, para. 12.

¹⁸For an interesting take on the brain in a vat hypothesis, see Nick Bostrom, “Are You Living in a Computer Simulation?” in *The Philosophical Quarterly* 53 (2003).

¹⁹An extremely helpful rundown of skeptical arguments and potential responses to them is to be found in Barry Stroud, *The Significance of Philosophical Skepticism* (Oxford: Oxford University Press, 1984).

assess its probability. We have no way of determining the truth or falsity, or even asserting the probability or improbability, of these various skeptical scenarios. *This is a necessary consequence of metaphysical realism's division of our best understanding of the world and the world itself*: after all, even with the best possible theory of the world, there still remains the possibility that this theory is incorrect: the skeptical scenarios simply vividly illustrate metaphysical realism's constitutive cleavage.

By this point, the story I'm telling should seem familiar. Metaphysical realism seems to generate the same epistemic circumstances with regard to claims about the external world that consequentialism generates about the moral valence of actions. Consequentialism faces a skeptical scenario: that the balancing-out hypothesis is false, i.e., that the invisible consequences of any two actions ϕ and ψ are sufficient to outweigh the visible consequences of ϕ and ψ in assessing the comparative moral valence of ϕ and ψ . For consequentialism, the probability of this skeptical scenario is beyond our epistemic grasp, given that invisible consequences are inaccessible to even the most epistemically responsible agent. The same is true of metaphysical realism. Whether the computer in front of me is a real computer or merely a simulation is inaccessible to me, as is the probability that it is real or merely the product of an evil demon's imagination.

5. *Three Responses*

Given the similar predicaments of metaphysical realism, as concerns our knowledge of the external world, and consequentialism, as concerns our knowledge of the moral valence of actions, it is helpful to assess the various strategies that have been employed in responding to the challenge of epistemological skepticism. As far as I can see, there are three broad methods of response. The first is to accept metaphysical realism and to deny skepticism. The standard, and most obvious, method of occupying this position is to put the blame for epistemological skepticism not on the existence of an external world, but rather on the requirements of knowledge or epistemic access. According to this option, we can have knowledge of the external world without access to the falsity of skeptical scenarios. The second option, less widely accepted, is to accept that metaphysical realism implies skepticism, but to accept metaphysical realism anyway. A standard view in this general area is to recognize that though skepticism is true, we can continue to build theories of the world, admitting that we cannot genuinely *know* whether these theories are true or false. The third option is to deny metaphysical realism. I discuss each of these possibilities in turn.

I should briefly state what I hope to be the upshot of my discussion in this section. I hope to show that those who would reject consequentialism on the basis of arguments from cluelessness face a dilemma. They must either reject metaphysical realism also (in effect, accepting the third option), or face an undischarged argumentative burden: that skeptical challenges cause problems for consequentialism but not metaphysical realism. By the end of §6, I hope to have shown that neither horn of this dilemma is attractive.

5.1. *The Epistemic Option*

Perhaps the most common response to epistemological skepticism is to challenge the various epistemological assumptions that require us to say that the possibility of skeptical scenarios renders propositions about the external world unknowable. Though I do not wish to provide an exhaustive catalogue here, some of the more popular skepticism-defeating²⁰ epistemic options run as follows:

1. *Fallibilism*: knowledge that p is possible even if one cannot rule out the possibility that $\neg p$.
2. *Contextualism*: “knowledge” is contextually defined. In so-called “everyday” contexts, one can have knowledge that p without ruling out the possibility that $\neg p$. In more stringent contexts, one must rule out the possibility that $\neg p$ to have knowledge of p .
3. *Denial of Epistemic Closure*: if one knows that p , and knows that p entails q , one need not know q .

Of course, each of these principles is only very roughly articulated and permits of a number of modifications and fine-tunings. Leaving this aside, take, first, fallibilism. According to Fantl and McGrath, fallibilism is motivated in the following way:

[I]magine a brain in a vat with the same experiential states and apparent memories as you have. It is at least somewhat plausible to think that this subject has just as good evidence as you have for many ‘obvious’ claims, such as *I have a body*, *I interact daily with a number of other people*, etc. But this subject is completely wrong.

²⁰I do not wish to engage in a discussion of whether any particular option *genuinely* defeats skepticism; rather, I seek only to show that each of these possibilities, if it works for metaphysical realism, also works for consequentialism. I needn’t accept the antecedent in all cases.

What should we make, epistemically, of these ‘bad cases’? One familiar sort of skeptic, of course, wants to appeal to the bad cases to show that, contrary to what we think, we *don’t* have knowledge in the corresponding ‘good cases’. Against this, the fallibilist holds her ground: the existence of these bad cases, and even our acknowledgement of them, does not undermine our knowledge in the good cases. The fallibilist allows that we can have *fallible* knowledge.²¹

Should this view succeed, a response to skepticism looms. I may know that I have two hands despite the fact that there is a proposition that is incompatible with this claim (e.g., that I am being deceived by an evil demon) that I cannot rule out. Why? Because any knowledge that I would have for any proposition is compatible with my fallibility; it is compatible with the fact that I cannot rule out being deceived by an evil demon.

Though whether fallibilism genuinely succeeds is up for dispute, *if* it succeeds, this is a success not only for metaphysical realism, but also for consequentialism. Because the claim that if I can be said to know p (the moral valence of ϕ) without positively ruling out some proposition q (the falsity of the balancing-out hypothesis) that is incompatible with p , it would appear I can know that my action of giving \$100 to Oxfam International is morally required, even though I cannot rule out the possibility that the invisible consequences of so doing will outweigh the comparative value of the visible consequences of not so doing.

A further method to reject skeptical scenarios is to restrict their applicability to certain *contexts*.²² In the most stringent, or “philosophical” context, knowledge of the external world requires evidence that one is not being deceived. Because there is no such evidence, the philosophical context accepts skepticism. But, on the street, at a party, making plans with friends, in the lab, etc., things are not so stringent. Claims of knowledge do not require us to rule out skeptical scenarios.

Again, it is not my task to address the plausibility of contextualism as a solution to epistemological skepticism. But it is clear that a contextualist epistemology might work just as well for consequentialism as for metaphysical realism. For instance, we might suggest that our inability to discern the invisible consequences of a particular action entails that we do not know the

²¹Jeremy Fantl and Matthew McGrath, *Knowledge in an Uncertain World* (Oxford: Oxford University Press, 2009), 6-7.

²²See, for instance, David Lewis, “Elusive Knowledge” in *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1998).

moral valence of a particular action only in certain, highly circumscribed, contexts. (Including, perhaps, the philosophical discussion of the epistemic merits of consequentialism.) But while on the street, at a party, making plans, educating children, designing decision procedures, etc., one needn't take seriously our inability to know the invisible consequences of a given action because, when it comes to knowledge-claims about the moral valence of actions, we are licensed to rule out the falsity of the balancing-out hypothesis and judge the moral valence of actions on visible consequences only, just as we are licensed to make knowledge claims of the external world based only on visible features of the external world, leaving aside skeptical scenarios.

Another epistemic option is to deny a standard feature of epistemological reasoning. Skepticism appears to get off the ground only if one accepts what might be called the *Epistemic Closure Principle*.²³ Roughly speaking, this principle holds that if I know p , and know that p entails q , I thereby know q . On the assumption of metaphysical realism, one might construct an argument for skepticism in the following way:

1. I know p (where p is an ordinary claim about the external world).
2. I know p entails $\neg q$ (where q is a skeptical scenario).
3. If I know p and I know that p entails $\neg q$, I know $\neg q$.
4. Hence, I know $\neg q$.
5. But I do not know $\neg q$.
6. Hence, I do not know p .

If one denies (3), i.e., the closure principle, one can avoid the problematic *reductio*, and claim both that one does not know that the skeptical scenario is false, but also that one does know the ordinary claim about the world, or about the moral valence of a particular action.

As with fallibilism and contextualism, I do not wish to comment on the plausibility of the denial of the epistemic closure principle. In any event, one should once again realize that *if* we accept this epistemic solution to the problem of skepticism in the case of metaphysical realism, there is no reason to bar this solution for the problem of cluelessness for consequentialism: I can know that assisting a drowning child was a morally better action than not having done so on the basis of visible consequences, without also knowing that invisible consequences do not alter the moral valence of the rescue.

²³For a good introduction to this principle and its status see, Luper, Steven, "The Epistemic Closure Principle", The Stanford Encyclopedia of Philosophy (Spring 2009 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2009/entries/closure-epistemic/>>.

In considering these possibilities, a general conclusion can be drawn. *Any epistemological principle that allows knowledge of the external world allows knowledge of the consequentialist moral valence of actions.* For consequentialism and metaphysical realism, we have epistemic access to that which is “visible” (i.e., visible consequences, perceptions, etc.). But if we accept an epistemology that allows knowledge of the external world simply on the basis of what is “visible”, this epistemology must allow knowledge of consequentialist moral evaluations on the basis of what is “visible”. There can be no epistemic distinction, insofar as the respective skeptical scenarios are on the same epistemic footing.

One response is possible here. Look again at epistemological fallibilism. Fallibilism claims that one can know p without having sufficient evidence to rule out all propositions that imply $\neg p$. But to know p , one must at least have some *threshold* of evidence in favor of p ; in slightly different terms, one must have *some* epistemic reason to believe p . As Fantl and McGrath put the matter: “If fallibilism of our strong epistemic sort is true, then there can be knowledge without probability 1. But presumably there can’t be knowledge with probability 0, probability 1/2, or even probability 2/3!... These considerations raise the question of just how probable a proposition must be for you in order for you to know it.”²⁴ A similar point applies when it comes to the denial of epistemic closure. One must be able to assert, relatively plausibly, that (as in (1)) “I know p ”. But “I know p ” plausibly requires some epistemic reason in favor of p . If so, we may have a way to distinguish metaphysical realism and consequentialism. For metaphysical realism, the claim that I know I have two hands, for instance, is itself *plausible*. Surely I have *enough* evidence for the claim that I have two hands to know it, even if I cannot rule out its falsity (given skeptical scenarios). But it is not similarly plausible to say that I know that my donation to Oxfam International is morally required. It would appear that I do not have sufficient evidence to make this claim.

This response fails. The consequentialist does have at least some reason to accept a particular verdict with regard to the moral valence of ϕ in comparison to ψ : the visible consequences of ϕ and ψ . Someone might claim that this is not enough. But on what basis? It seems that *the only* basis to conclude that this is insufficient evidence is the potential falsity of the balancing-out hypothesis. But if this is sufficient to render one’s knowledge of the moral valence of ϕ void, it would appear that the inability to rule out the falsity of skeptical scenarios is also sufficient to render one’s knowledge

²⁴Fantl and McGrath, 25.

that one has two hands void. After all, the only epistemic reason not to believe the evidence of one's bi-handedness (i.e., its perceptual appearance) is be the possibility that, e.g., one is in a computer simulation (or is being controlled by an evil demon, etc., etc.). In both cases, available evidence is enough to warrant knowledge of the particular proposition *but for* the *invisible* skeptical scenario. One might say that the evidence that I have two hands is comparatively weightier than the evidence I have that ϕ is morally justified on a consequentialist regime. But one can only make this comparative claim if one has epistemic access to the evidence in favor of skeptical scenarios, which one does not have, or that one can be justified in believing them false, which one cannot be. Hence, I can see no reason to believe that there is less weighty evidence for the moral valence of actions under consequentialism than any particular claims about the external world according to metaphysical realism. The moral valence of ϕ , and the claim that one has two hands, are perfectly equivalent from the point of view of a fallibilist epistemology, or any other epistemology that allows metaphysical realism to escape epistemological skepticism.²⁵

The key here is that for both metaphysical realism and consequentialism the skeptical scenarios are totally invisible. And if this is correct, there is no ground on which one might claim that the skeptical scenario is more or less plausible in one or the other case. To claim that beliefs about the external world are somehow on sounder epistemic footing than beliefs about the moral valence of actions on a consequentialist view is to assign epistemic reasons, or epistemic valence more generally, based on mere *whim*.

5.2. *Accept Skepticism, Soldier On*

One important possibility with which some have toyed is simply to accept metaphysical realism *and* its skeptical consequences. Perhaps, after all, we fail to have knowledge about the external world. In principle, one could adopt two courses of action under a skeptical regime. First, one could simply refrain from making any claims about the external world at all; one could simply admit one's ignorance in a *Pyrrhonian* fashion.

²⁵Consider Fantl and McGrath's own response to this question: "Your probability for p is knowledge-level iff the probability that not- p doesn't stand in the way of p 's being put to work as a basis for belief and action," (Fantl and McGrath, 26). But if this is correct, the probability that ϕ is morally permitted, given its visible consequences, surely passes the threshold of fallible knowledge: the skeptical scenario in this case certainly does not stand in the way of it being "put to work as a basis for belief and action," assuming that the consequentialist response to the *traditional* epistemic objection succeeds.

This reaction to skepticism is not particularly attractive, though it can be embraced by consequentialism. Consequentialism would require a simple refusal to make moral judgments, just as the Pyrrhonian will refuse to make judgments about the external world. This is compatible with the truth of consequentialism, just as Pyrrhonian skepticism is compatible with the truth of metaphysical realism. However, the most popular response to the fact of epistemological skepticism is to simply ignore it. When reflecting on his skeptical conclusions, Hume, for instance, writes:

Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium, either by relaxing this bent of mind, or by some avocation, and lively impression of my senses, which obliterate all these chimeras. I dine, I play a game of back-gammon, I converse, and am merry with my friends; and when after three or four hour's amusement, I wou'd return to these speculations, they appear so cold, and strain'd, and ridiculous, that I cannot find in my heart to enter them any farther.²⁶

Hume here appears to be suggesting that though skepticism is true, we simply ignore it: after real life takes over, and we have left our philosophical armchair, we forget about the possibility of skepticism and go on making claims about science, morality, backgammon, etc. In other words, *we soldier on*, doing what we can in the face of our failure of knowledge. We seek to gain something less than genuine knowledge, but significant enough to navigate the world and to make the sort of judgments we are inclined to make.

Quine seems to endorse something like this approach in “Epistemology Naturalized”. There Quine admits that “[t]he Humean predicament is the human predicament;”²⁷ skepticism, of the form I discuss here, is true. Nonetheless, this does not mean that we cannot go about the process of theory formation, or the process of information gathering just as we always have. We just admit that such processes cannot guarantee us knowledge in the classic sense. Quine writes: “The Cartesian quest for certainty had been the remote motivation of epistemology, both on its conceptual and its doctrinal side; but that quest was seen as a lost cause.”²⁸ In response, Quine

²⁶David Hume, *A Treatise of Human Nature* (Oxford: Oxford University Press, 2007 [1740]), 1.4.7.9. See also “An Abstract”, para. 27.

²⁷W. V. Quine, “Epistemology Naturalized” in *Ontological Relativity and Other Essays* (New York, NY: Columbia University Press, 1969), 72.

²⁸Quine, 74.

offers the following epistemological programme:

The stimulation of his sensory receptors is all the evidence anybody has to go on, ultimately in arriving at his picture of the world. Why not just see how this construction proceeds? Why not settle for psychology? Such a surrender of the epistemological burden to psychology is a move that was disallowed in earlier times as circular reasoning. If the epistemologist's goal is validation of the grounds of empirical science, he defeats his purpose by using psychology or other empirical science in the validation. However, such scruples against circularity have little point once we have stopped dreaming of deducing science from observations.²⁹

As I read Quine, his recommendation runs as follows. We can never really have the epistemic access to the external world that skeptics take themselves to be arguing against. In this sense, skepticism is true. But once we have stopped dreaming of such access, we are, in a sense, liberated. We work *within* that to which we *do* have access: our actual empirical science, which proceeds on the basis of the way the world, as a matter of fact, appears to us. Whether this empirical science actually represents the world as it really is is a question the answer to which cannot be known. But so what? We soldier on with the theories we have and the inquiries we can conduct within the limits of our imperfect sensory experience.

For Quine, the empirical scientist understands the limits of scientific inquiry, but nevertheless proceeds to develop a theory based on the evidence to which the empirical scientist has access. But the same procedure is available to the consequentialist. Here we accept that we are clueless when it comes to the overall consequences of any given action. Literally, we cannot know the *true* moral valence of an action. But the means by which we gather information about the consequences of our actions—surveying their *visible* consequences—remains. (Of course, this will require satisfactory responses to more traditional epistemic arguments, which—I stress—I am not investigating here.) Against the skeptical challenge, we soldier on from *within* the procedures to which we have access, leaving aside skeptical worries about that procedure itself. We gather what data we can about the future and past consequences of our actions, we deliberate in the *ex post* about such actions and accept and refine decision procedures (“theories”), based on these facts. Of course, we can never know for sure what the moral valence of any given

²⁹Quine, 75-6.

action is, any more than we can ever know for sure whether there really is an electron here, or whether these two liquids really did dissolve, or whether it is an elaborate hoax cooked up by an evil demon. But we make do with what we have. And that is good enough.³⁰

5.3. *Reject Realism*

The last alternative is perhaps the most radical. Some, when confronted with the possibility of epistemological skepticism, have decided to reject metaphysical realism altogether and accept some version of anti-realism: that facts about the world depend on, or are somehow settled, by our theories, actual or possible conceptual schemes, beliefs, etc.

Like the previous responses to skepticism, I shall not argue in favor or against the rejection of metaphysical realism as a response to skepticism. But consider where we find ourselves. To accept metaphysical realism, one can either reject or accept skepticism. If one rejects skepticism, one must adopt an epistemology that can also vindicate epistemic access to the moral valence of actions on a consequentialist moral theory. If one accepts skepticism, one either adopts a *Pyrrhonian* suspension of judgment, or decides to soldier on in the face of a lack of knowledge. Whatever option one adopts for metaphysical realism, there seems to be very little reason to deny that option for consequentialism in the face of the argument from cluelessness. If that's right, then it would appear that the only method by which one might respond to the challenge of epistemological skepticism that *does not* yield a similar response to the argument from cluelessness is to reject metaphysical realism. I take this to be an unattractive commitment of those who would reject consequentialism on grounds of cluelessness.

Consequentialism may very well be false. Consequentialism may be an assault on the integrity of the moral agent.³¹ Consequentialism may reject plausible claims about the nature of rights, or may railroad the distinction between doing and allowing, or intending harm and merely foreseeing harm. One can reject consequentialism for these reasons and more and salvage metaphysical realism. But to sensibly reject consequentialism on the basis of the argument from cluelessness, it would appear that one must deny

³⁰The "soldiering on" response to the problem of cluelessness is plausibly represented in Alistair Norcross, "Consequentialism and the Unforeseeable Future" in *Analysis* 50 (1990).

³¹Cf. Lenman, 367-370. See also Bernard Williams, "A Critique of Utilitarianism" in Smart and Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1974).

that any of the above possibilities can be applied to consequentialism. But why should this be? It would appear that the skeptical challenge is no less serious for metaphysical realism than it is for consequentialism. In both cases, the theories in question face a skeptical scenario (in the case of metaphysical realism, deceiver hypotheses, brain-in-vat hypotheses, etc.; in the case of consequentialism, the falsity of the balancing-out hypothesis) that is wholly invisible, for which there is no good reason to believe or reject. One is tempted, then, to claim that anyone who rejects consequentialism on the basis of cluelessness must insist that the above methods of responding to skeptical challenges are null and void, rendering metaphysical realism susceptible to epistemological skepticism in just the same way.

I don't want to make the strong claim, right now anyway, that anyone who rejects consequentialism on grounds of cluelessness must also commit to the rejection of metaphysical realism. However, I do wish to make a slightly weaker claim. There appear to be two options for those who would reject consequentialism on grounds of cluelessness. First, they can simply admit that skepticism also defeats metaphysical realism. Though I have no argument against those who would accept such a conclusion, it seems to me sufficiently problematic to warrant avoidance if at all possible. (In particular, it seems to me that if the choice is between admitting that an argument against consequentialism fails and rejecting metaphysical realism, the latter is a much higher cost to pay.) Second, they can attempt to argue that skeptical challenges cause serious problems for consequentialism in a way that they *do not* for metaphysical realism. But, as I have indicated above, and shall argue in the following section, this is a very difficult argumentative burden to discharge.

6. *Disanalogies*

To avoid the rejection of metaphysical realism, those who would reject consequentialism on the basis of cluelessness *had better argue* for the claim that skeptical scenarios are far more problematic for consequentialism than for metaphysical realism. I see two possible arguments for this claim. First, one might claim that the skeptical scenarios that haunt metaphysical realism and consequentialism are *not*, after all, on the same epistemic footing. We have more reason to accept that the balancing-out hypothesis is false than we have to accept the claim that we are brains in vats. Second, one might accept that the skeptical scenarios are on the same epistemic footing, but that skepticism is problematic for consequentialism in a way it is not for metaphysical realism: consequentialism is an *action-guiding* view, and

hence a lack of knowledge about the moral valence of our actions is far more debilitating than a lack of knowledge about the external world. Neither of these possibilities, it seems to me, draws a relevant distinction between consequentialism and metaphysical realism.

6.1. *Credibility*

One natural way to divorce the fate of consequentialism from the fate of metaphysical realism might be as follows. The skeptical scenarios that infect metaphysical realism and consequentialism are not equally powerful. The skeptical scenario offered by those who would accept the argument from cluelessness is, in a certain sense, perfectly imaginable. All it would take for any given act's moral valence to be plunged into uncertainty is for the invisible consequences of ϕ and ψ to overrule their visible consequences in an account of their comparative axiological value. But surely, as Lenman is quick to note, every action's consequences reverberate indefinitely. Hence it is perfectly imaginable that the skeptical scenario, in the case of consequentialism, holds.

The skeptical scenario that plagues metaphysical realism, however, seems not a little outlandish. The skeptic would have us believe that it is possible that we are *brains in vats*, or that our perceptions are being controlled by an *evil demon*. But these possibilities are beyond the realm of believability. This response is not meant to show that there is a difference *in kind* between skeptical arguments against metaphysical realism and skeptical arguments against consequentialism, merely degree. One might put the point this way: for consequentialism, there seems to be substantially stronger epistemic reason to believe in the skeptical scenario than for metaphysical realism. If so, it would appear that, e.g., a fallibilist epistemology can vindicate metaphysical realism and not consequentialism: there is far more *epistemic reason* to believe the skeptical scenario in the case of consequentialism than there is in the case of metaphysical realism.

I have two responses. First, this response appears to let “seemings” determine the epistemic valence of actions. But in so doing, this response simply capitulates to Cowen's defense of consequentialism's epistemic access to “big” actions. For those actions, the balancing-out hypothesis *seems* true. Furthermore, as I argued above, it *seems* plausible to accept the balancing-out hypothesis (given the sheer number of consequences of actions with invisible consequences) holds even for comparatively small actions: it seems quite difficult to believe that my donation of \$100 to Oxfam International could be morally wrong simply on the basis of some unknowable invisible

consequences.

Of course, that it *seems* plausible to me that the balancing-out hypothesis holds doesn't say anything about whether it *seems* plausible to anyone else. Surely our intuitions on this score can and will diverge. But this leads to the second response to this argument. As Lenman himself is at pains to establish, that *p seems* correct has no status in determining the epistemic valence of *p*. Even if it *seems* to us that the skeptical scenario in the case of metaphysical realism is outlandish, or that there is less epistemological reason to believe it than to believe in the falsity of the balancing-out hypothesis, we have no *basis* for making this claim. To say that the falsity of the balancing-out hypothesis is less outlandish than, e.g., that we are brains in a vat requires us to have at least some epistemic access to invisible consequences. Insofar as we have no such access, however, we have no business claiming that there is or is not stronger or weaker epistemic reason to accept the skeptical scenario. The claim that *p* "seems outlandish" offers no justification for believing that *p* is false. Of course, if we had some bit of evidence that favored the rejection of the balancing-out hypothesis, we might plausibly claim that there is stronger epistemic reason to believe it than there is to believe that we are brains in vats. But we could only come to have this evidence if we have epistemic access to invisible consequences, a *lack* of access to which the argument from cluelessness relies upon. Hence the fact, if it is a fact, that skeptical scenarios faced by metaphysical realism are outlandish where the skeptical scenarios faced by consequentialism are not should be given no argumentative weight at all. Insofar as epistemic reasons are based on available evidence, the epistemic reason to believe the falsity of the balancing-out hypothesis is of precisely the same weight as the epistemic reason to believe that we are brains in vats. We have no evidence for or against either claim.

One response to this argument should be considered here.³² We know, of course, that many good and bad things have happened. Take, for instance, the Emancipation Proclamation. But we also know that the end of slavery in the United States was an invisible consequence of some action or other (though, of course, we cannot know which one given that it is an *invisible* consequence). And if this is the case, why don't we have some evidence that, at least in some cases, invisible consequences will far outweigh visible consequences? This response does not succeed. To have evidence that invisible consequences outweigh visible consequences, one must not simply know that good things have happened, and that these things are invisible consequences

³²Thanks to an anonymous reviewer.

of some action or other. Rather, one has to have some epistemic access to the *balancing-out hypothesis*: whether, for any *particular* act, the invisible consequences of *that* action balance out. But we have no such access—as the argument from cluelessness states explicitly—insofar as we have no capacity to assign such invisible good and bads to particular actions (given that they are invisible). To assess the balancing-out hypothesis for any action, we are led back to epistemically impotent “seemings”.

6.2. *Morality is Special*

There is no reason to believe that any epistemological fix for metaphysical realism couldn't also allow consequentialism to avoid the spectre of the falsity of the balancing-out hypothesis. But this doesn't mean that the similar epistemic footing of skeptical scenarios are similarly challenging to consequentialism and to metaphysical realism. Consequentialism, as a moral theory, *is action guiding*. Hence if we have no knowledge of the moral valence of our actions, there is reason to be worried: when we don't know, we can't guide action. In the case of metaphysical realism, what we don't know won't hurt us; metaphysical realism is a theory about the nature of the external world, and hence is not an action guide. Skepticism is *much* more problematic, one might claim, in the case of an action-guiding view than in the case of a non-action-guiding view like metaphysical realism. Moral skepticism is debilitating in a way that garden-variety skepticism is not.

There are two problems with this response. First, it succeeds only if the proper response to skepticism in any domain is *Pyrrhonian*. If we allow that we can *soldier on*, or revise our epistemological principles, skepticism is no more problematic for consequentialism than for metaphysical realism. If we accept either alternative, we can guide action by focusing on the visible consequences, and ignoring the possibility that the balancing-out hypothesis is false. Second, and more seriously, there can be no plausible distinction between consequentialism and metaphysical realism on grounds that skepticism about the latter, but not skepticism about the former, leaves our actions susceptible to guidance. Skepticism about the external world *implies* skepticism of the moral valence of actions on a very wide range of moral theories.³³ Take, for instance, an explicitly non-consequentialist theory, such as a *rights*-based view. Now let's say that Jerry murders Frank.

³³The following response doesn't hold if this moral theory never treats any proposition about the external world as a moral reason. Such views are possible, but seem so extreme as to render Lenman's so-called “disengaged consequentialism” (Lenman 361-2) positively plausible in comparison.

It is certainly true that no apparitions conjured up by evil demons have rights. No simulations used by computers to manipulate our brains have rights either. Is Jerry's murder of Frank morally permissible? If Jerry is a brain in a vat, yes: he destroyed a non-rights-bearing entity, i.e., a computer simulation of another person. In this case, Jerry's action is no more morally blameworthy than the act of playing *Grand Theft Auto III*. If Jerry is not a brain in a vat, or is not being deceived by an evil demon, the answer is no: Jerry *actually* killed Frank, who possesses moral rights not to be killed. Hence, on *any* moral theory that accepts that propositions about the external world can be moral reasons, the truth of metaphysical realism affects our ability to guide action. Hence we should be no more or less sanguine about skepticism when it comes to metaphysical realism as we are about skepticism when it comes to the invisible consequences of our actions.

It seems to me that there is no good reason to believe that any metaphysical realism-preserving stance one might take toward epistemological skepticism does not allow consequentialism to fend off the argument from cluelessness. Of course, there are many strong arguments against consequentialism. Indeed, it may be that the *traditional* epistemic argument against consequentialism is reason enough to reject it. But the argument from *cluelessness*, as distinct from the traditional argument, proves too much.

7. Conclusion

Those who would reject consequentialism on basis of cluelessness face a dilemma. Either they must reject metaphysical realism, or they must show why methods to salvage metaphysical realism in the face of skepticism are not also available to those interested in salvaging consequentialism. But insofar as there appears to be no principled reason to embrace the latter horn, this is a sticky situation for the proponent of the argument for cluelessness.

One question might arise. How do we conduct consequentialist analyses of actions under the threat of epistemological skepticism? The answer is easy, and has been hinted at already. Moral inquiry will proceed precisely as Mill recommends: by judging the moral quality of actions on the basis of their *visible* consequences. Guidance of action will occur—as it should—by considering the morality of actions in the *ex post*, and by constructing secondary rules and principles based on these judgments. We may well admit, in our reflective moods, that we do not really *know* that the visible consequences determine the moral valence of actions. (Or, perhaps, that our knowledge is fallible, or that we cannot rule out the possibility of the falsity of the balancing-out hypothesis.) But we will make judgments and

guide actions, just as those who admit the possibility of skeptical scenarios persist in a belief in their best scientific theories.

Before I conclude, I should note what I have not done in this paper. For starters, I have defended consequentialism only from the argument from cluelessness. I have not defended consequentialism from the more traditional epistemic objections. Second, I have no argument against someone who wishes to accept that consequentialism fails in light of epistemic cluelessness *and* that metaphysical realism fails in the face of epistemological skepticism. Though I take this to be an unattractive position, it is open. Furthermore, there may be some argument I have not foreseen that shows the special problem of skeptical scenarios for consequentialism in contrast to metaphysical realism. Though I doubt such an argument could succeed, I make no *general* claim about its impossibility. I do, however, wish to insist that those who would reject consequentialism on epistemic grounds face an undischarged argumentative obligation: they must draw this distinction plausibly. Though I am skeptical, I leave open the possibility that this could be done.