

The Infectiousness of Nihilism*

William MacAskill

In "Rejecting Ethical Deflationism," Jacob Ross argues that a rational decision maker is permitted, for the purposes of practical reasoning, to assume that nihilism is false. I argue that Ross's argument fails because the principle he relies on conflicts with more plausible principles of rationality and leads to preference cycles. I then show how the infectiousness of nihilism, and of incomparability more generally, poses a serious problem for the larger project of attempting to incorporate moral uncertainty into expected value maximization style reasoning.

In "Rejecting Ethical Deflationism," Jacob Ross argues that a rational decision maker is permitted to reject certain moral theories: that is, assume that such theories are false for the purposes of practical reasoning.¹ Consider, for example, a "uniform" theory, according to which all options are equally choice-worthy, and suppose that a decision maker is less than certain in such a theory. Every time that such a decision maker encounters a choice-situation, the portion of her credence devoted to this uniform theory will give no reasons in favor of any option, whereas the portion of her credence in positive moral theories will give reasons in favor of some options over others. In every choice-situation, therefore, the same actions will be rationally permissible and impermissible as if she had zero credence in the uniform theory. So a rational decision maker need never worry about the portion of her credence devoted to this uniform theory.

* For very helpful comments on this article, I am indebted to Nick Beckstead, John Broome, Krister Bykvist, Amanda MacAskill, Toby Ord, Jacob Ross, the Editor at *Ethics*, and the audience at New York University's Washington Square Circle. Parts of this essay were written while I was the recipient of a UK-US Fulbright Scholarship and the recipient of an Arts and Humanities Research Council Award.

1. Jacob Ross, "Rejecting Ethical Deflationism," *Ethics* 116 (2006): 742–68.

Ethics 123 (April 2013): 508–520

© 2013 by The University of Chicago. All rights reserved. 0014-1704/2013/12303-0001\$10.00

What's striking about Ross's conclusion is that it holds no matter how confident you are in that uniform theory, as long as that confidence does not amount to full certainty. Even if you are 99.9 percent certain that the uniform theory is correct, you should still reject that theory and act in the way determined by the remaining 0.1 percent of your credence. This goes against the implicit received view of decision making under moral uncertainty—known as the “my favorite theory” account²—according to which one should simply follow the theory one thinks is most likely to be correct.

Ross extends his argument to other theories. He discusses certain sorts of relativistic theories, and certain “relatively deflationary” theories, and argues that they can be rejected too. I will not focus on these arguments. Rather, I will focus on his argument that one may rationally reject nihilism. Insofar as nihilism is a widely held view, Ross's argument that we can rationally reject nihilism seems particularly important.

In this article, I hope to demonstrate two things. First, I hope to demonstrate that Ross's argument for the conclusion that one can rationally reject nihilism fails. I do not claim that it's *false* that one can rationally reject nihilism. I only claim that Ross has not given us solid grounds for thinking that we should rationally reject nihilism.

Second, I hope to demonstrate that the fact that one should have nonzero credence in nihilism—or, more generally, in theories that posit extensive incomparability³—poses a serious problem for a larger project of which Ross's article is merely a part. This larger project is the attempt to extend expected value maximization style reasoning to encompass moral uncertainty as well as descriptive uncertainty.⁴ Again, I am not claiming that attempts to apply expected value maximization style reasoning under moral uncertainty are doomed to failure. I merely claim that what I will call the problem of ‘infectious incomparability’ is a powerful challenge that has not yet been raised but that needs to be met.

The structure of my argument is as follows. In Section I, I define nihilism and show why ‘expected value’ reasoning can't be used to reject nihilism; this also serves to introduce the infectiousness problem. In Section II, I consider Ross's dominance principle, which he uses to argue that

2. Ted Lockhart, *Moral Uncertainty and Its Consequences* (Oxford: Oxford University Press, 2000), 42.

3. By “*x* is incomparable with *y*” I simply mean that no positive value relation holds between *x* and *y*. See Ruth Chang, “Introduction,” in *Incommensurability, Incomparability, and Practical Reason* (Cambridge, MA: Harvard University Press, 1997), 4.

4. This project has been developed in Lockhart, *Moral Uncertainty and Its Consequences*; and Andrew Sepielli, “Review of Ted Lockhart's ‘Moral Uncertainty and Its Consequences,’” *Ethics* 116 (2006): 601–4, “What to Do When You Don't Know What to Do,” in *Oxford Studies in Metaethics*, vol. 4, ed. Russ Shafer-Landau (Oxford: Oxford University Press, 2009), 350, and “Moral Uncertainty and the Principle of Equity among Moral Theories,” *Philosophy and Phenomenological Research* (2012), doi:10.1111/j.1933-1592.2011.00554.x.

we can rationally reject nihilism. I give two arguments to show why Ross's principle is false. In Section III, I consider two responses that Ross has made in correspondence and argue that they fail. In Section IV, I consider a rejoinder but argue that that rejoinder fails. I conclude that the problem of infectious incomparability is a grave problem for attempts such as Ross's to incorporate moral uncertainty into expected value reasoning.

I

Ross defines nihilism as "the view that the notions of good and bad and of right and wrong are illusions and that, objectively speaking, no option or state of affairs is better than any other, nor are any two options or states of affairs equally good. Thus, while uniform theories assign the same value to all of our options, nihilistic theories don't assign values to any of our options."⁵ Notice that nihilism, on this definition, is very different from the uniform theory mentioned above. According to the uniform theory, a positive value relation obtains between all options: every option is equally as good as every other option. According to nihilism, *no* positive value relation obtains between any two options. That is, the value of every option is undefined.

However, if the value of every option is undefined, according to nihilism, then, for any decision maker with nonzero credence in nihilism, there is a big problem for her if she attempts to incorporate moral uncertainty into her reasoning about expected value. The problem is as follows. If we take an expectation over possible states of nature, taking the sum, for each state of nature, of the value of that state of nature multiplied by its probability, and the value of one state of nature is undefined, then the expectation as a whole is undefined. Because, according to nihilism, the value of every option is undefined, for a decision maker with nonzero credence in nihilism, the expected value of every option is undefined, too. Nonzero credence in nihilism is therefore sufficient to infect practical reason, resulting in there being no subjective reason for preferring any option over any other.

This is pretty terrifying. Nihilism isn't just a speculative hypothesis. It is a view that's defended by many intelligent philosophers. It's difficult to see how it could be logically false, or somehow conceptually confused. So it seems exceedingly plausible that we ought to have nonzero credence in the view. And it is also very plausible that we should take moral uncertainty into account in our reasoning about the expected value of different paths of action. But if we grant all this, then, for each and every one of us, the expected value of all options is undefined.

5. Ross, "Rejecting Ethical Deflationism," 748.

II

Ross, however, does not use expected value reasoning in order to argue that one ought to reject nihilism. Rather, he uses dominance reasoning. He supposes, as an illustration, that he has credence in two theories only: “ T_L ,” which is a positive moral view, and “ T_n ,” a nihilist theory. He argues:⁶

According to T_L , it would be better for me to send the trolley to the left than to send it to the right. And so my credence in T_L gives me *pro tanto* subjective reason to send the trolley to the left. The only way this could fail to be the most rational option would be if my credence in T_n gave me a sufficiently strong countervailing subjective reason to send the trolley to the right. But T_n implies that there would be nothing valuable or disvaluable about either alternative. And so my credence in T_n gives me no subjective reason to favor either alternative. Hence the *pro tanto* subjective reason to send the trolley to the left is unopposed, and so this is the rational option.

That is, Ross appeals to something like the following principle:

Dominance over Theories (DoT).—If some theories in which you have credence give you subjective reason to choose x over y , and no theories in which you have credence give you subjective reason to choose y over x , then, rationally, you should choose x over y .⁷

Note that DoT is not a strong enough principle to warrant Ross’s conclusion that one can always reject nihilism. It merely warrants the rejection of nihilism in cases where one has credence in only one theory other than nihilism, and only a small number of cases in which one has credence in more than one nonnihilist theory. We could easily formulate a stronger principle, which would give Ross the result he wants. But I’ll focus on this weaker principle for ease of exposition; all the same considerations I raise would apply *mutatis mutandis* to the stronger principle. I will give two arguments to show why DoT is false.⁸

Argument 1. DoT Conflicts with More Plausible Principles of Rationality

Suppose, for example, that the decision maker has credence in the following two theories only: first, a specific person-affecting view of population ethics. The basic idea behind this view is that you can only have

6. Ibid.

7. Ross endorsed a very similar principle in correspondence, but then after discussion suggested a modified principle, which is discussed in Sec. III.

8. These arguments assume that rational choice should not be cyclical. This assumption is discussed in Sec. III.

reasons to make people better off (or to avoid making people worse off), and that one cannot be made better off by being brought into existence, so there is no reason simpliciter to bring someone into existence, even if that person would have a happy life. According to this view, population *X* is better than population *Y* iff the two populations have exactly the same people in them and there is a greater average well-being in population *X* than in population *Y*. If the populations do not have exactly the same people in them, then it is undefined whether *X* is better than *Y* (or, in other words, *X* is incomparable in value to *Y*). If *X* and *Y* are entirely distinct populations, then there is not even a *pro tanto* reason to bring about *X* over *Y*, or vice versa, because there's no one for whom *X* is better than *Y*, or for whom *Y* is better than *X*.

The second theory in which the decision maker has credence is egalitarian utilitarianism, which places value both on average well-being and on the equality of the distribution of that well-being, and according to which these two values can be weighed against each other. The theory gives a complete ordering of populations in terms of choice-worthiness.

Now consider three populations one could bring about: populations *A*, *B*, and *C*. *A* and *C* have the same people in them. *C* has a much greater average well-being than *A*, but the well-being in *A* is distributed much more equally. Population *B* is larger in size than populations *A* and *C*, and no person in population *B* exists in population *A* or *C*. Population *B* has the same average well-being as *A*, but a slightly less egalitarian distribution of well-being. According to egalitarian utilitarianism, *A* is better than *B*, which is better than *C* (i.e., $A > B > C$). According to the person-affecting view, *A* and *B* are incomparable, *B* and *C* are incomparable, and *C* is better than *A* (i.e., $A \sim B$; $B \sim C$; $C > A$).

Now let's suppose the decision maker finds the person-affecting view much more likely: she has 99 percent credence in the person-affecting view, and only 1 percent credence in egalitarian utilitarianism. And let's suppose that there is a huge amount at stake, in the choice between *A* and *C*, according to the person-affecting view, whereas not much is at stake in the choice between *A* and *C*, according to the egalitarian utilitarian view. Let's suppose that the value difference between *C* and *A*, according to the person-affecting view, is 100 times greater than the value difference between *A* and *C* according to the egalitarian utilitarian view. Doing the math, the expected value of *A* is $0.01 * 1 = 0.01$ and the expected value of *C* is $0.99 * 100 = 99$.

Intuitively, the decision maker should certainly prefer *C* to *A*: the situation is both much higher stakes according to the person-affecting view, and the decision maker has much higher credence in the person-affecting view. But, by DoT, the decision maker should prefer *A* to *B*, because egalitarian utilitarianism gives some reason in favor of preferring *B* to *A*, and the person-affecting view gives no reason in favor of pre-

ferring *B* to *A*. And, by exactly similar reasoning, DoT implies that the decision maker should prefer *B* to *C*. Therefore, assuming acyclicity, she should not also prefer *C* to *A*. So DoT conflicts with our intuitive appraisal of the above situation.

In the above case, both the size of the credence in the person-affecting view and the magnitude of the ratio of the value difference between *A* and *C* across the two theories were arbitrary. So let's take that into account, and generalize the above objection into a principle, which we can call the weak comparativism principle (WCP):

WCP.—If one has credence p in T_1 , and credence $(1 - p)$ in T_2 (where $1 > p > 0$), and x is more choice-worthy than y according to T_1 , and y is more choice-worthy than x according to T_2 , and the difference in choice-worthiness between x and y is n times greater, according to T_1 , than it is according to T_2 , then, for some p and n , one should choose x over y .

WCP is in conflict with DoT. But WCP is pretty darn plausible: it is about the weakest nondominance principle that I can think of. It is much weaker, for example, than the principle that one should prefer x over y if x has higher expected value than y . It seems much more plausible than DoT.⁹ So we should reject DoT.

Argument 2. DoT Generates Preference Cycles

To see this, consider the following case. Suppose that you are the leader of the World Government, and you are to decide what policies to pursue. The principal questions facing you are whether to pursue a policy of population control, and to what extent you should mine natural resources, improving technological and economic progress but at the cost of environmental preservation.

You have many options available to you, but here are three. Option *A* involves population control and moderate use of scarce resources. This will result in 1 billion people at well-being level 4, 1 billion people at well-being level 2, and moderate environmental destruction. Option *B* involves population control, but the preservation of scarce resources. This will result in 2 billion people all living at well-being level 2, with no environmental destruction. Option *C* involves no population control and the preservation of scarce resources. This will result in 5 billion people living at well-being level 2, and 5 billion people living at well-being level 1 with no environmental destruction.

9. This seems especially true when we consider, as we will in argument 3, that we can formulate a principle that is very similar to DoT, which does not suffer from the same conflict.

You are uncertain about what the true moral theory is. You split your confidence between three theories only. The first, T_1 , is a person-affecting view. According to this view, when comparing two distributions of well-being between populations that have the same number of people, one should choose the distribution with the highest average well-being. Populations involving different numbers of people, on the other hand, are absolutely incomparable in value. On this view, environmental preservation is of no value. So, according to this theory, A is better than B , and C is incomparable with A and B (i.e., $A > B$; $B \sim C$; $A \sim C$).

The second, T_2 , is an environmentalist theory. On this theory, average well-being and environmental preservation are both of value, so one should choose policies that increase average well-being and choose policies that preserve the environment. But, on this view, average well-being and environmental preservation are absolutely incomparable in value: there are no facts about how one should weigh average well-being and environmental preservation. So, according to this theory, B is better than C , and A is incomparable with B and C ($B > A$; $A \sim C$; $A \sim B$).

The third theory, T_3 , is an egalitarian theory. It values both the sum total of well-being and the level of equality of well-being among people. But, on this theory, the value of total well-being and the value of equality are absolutely incomparable: there are no facts about how one should weigh one against the other. So, according to this theory, C is better than A , and B is incomparable with C and A (i.e., $C > A$; $B \sim C$; $B \sim A$).

In the above case, DoT entails that one ought to have cyclical preferences. The above example is a bit complex to think through, so I lay out the theories in the following tabulation, which makes things easier to see. Each theory has two columns underneath it. If one option is above another in the same column, then one has greater reason to choose the higher option than the lower option. If one option is in a different column than the other option, then the two options are incomparable.

T_1		T_2		T_3	
A	C	A	B	C	B
B			C	A	

According to DoT, you ought to prefer A over B , B over C , and C over A . That is, DoT entails that sometimes we ought rationally to have cyclical preferences among the options in our choice-set.

III

In response to these arguments, one might be tempted by one or both of two responses. In this section, I argue that these responses fail.

A. *Modify DoT*

First, one might be tempted to modify DoT. In correspondence, Ross has suggested the following modified principle in place of the principle suggested in his article:

Modified Dominance over Theories (MDoT).—For any action φ , rationality requires that if, among the theories in which you have positive credence, some imply that you have most objective reason to φ , and none imply that you have any objective reason not to φ , then you intend to φ .

The key change to this principle is that, while DoT appealed to the idea that one's credence in a theory might not give one subjective reason overall for preferring x to y over vice versa, MDoT appeals to the idea that a theory might entail that there are *no* reasons against choosing a particular option.¹⁰ To illustrate the importance of this distinction, consider that there are two ways in which incomparability can arise. First, it can arise because there are reasons in favor of one option, and reasons in favor of another option, but that those competing reasons cannot be weighed against one another, and so the two options are incomparable. A theory might, for example, claim that there is some reason to become a lawyer (because becoming a lawyer will benefit others), and some reason to become a clarinettist (because becoming a clarinettist furthers one's aesthetic ideals), but that these reasons cannot be weighed against each other—in which case, having credence in that theory would not give one subjective reason for preferring to become a lawyer over a clarinettist or vice versa. But it would still be the case that there is some objective reason, according to that theory, not to become a lawyer, and some objective reason, according to that theory, not to become a clarinettist.

Second, incomparability can arise because some precondition of two options' being comparable is not met. For example, if nihilism is true then, though becoming a lawyer and becoming a clarinettist are incomparable, that is not because there are reasons in favor of each, and those reasons cannot be weighed against each other. Rather, the options are incomparable, according to nihilism, because there are no objective reasons at all. So credence in nihilism gives one *no* subjective reason to prefer becoming a lawyer over becoming a clarinettist or vice versa. But, unlike in the pre-

10. Two other changes—that the principle is now explicitly wide scope, and that we are now talking about rational intention rather than rational choice—are not relevant to the arguments here discussed. The fact that the principle now only refers to an option's having most reason in favor of it, rather than being a comparative relation between two options, means that the cyclicity generated is over choice-sets, rather than within a choice-set. Given the arguments that follow, this also isn't an issue. Though these changes distract from the main issue somewhat, I wished to be true to Ross's new formulation.

vious case, this is because there's no objective reason against becoming a lawyer, and no objective reason against becoming a clarinetist.

Most instances of incomparability are similar to the former example rather than the latter, and so MDoT, while allowing one to reject nihilism, will run into trouble less often than DoT because the kind of dominance it defines kicks in only in when the dominated theories offer no reason at all pertaining to the options in question. But it still gets into trouble. Consider the example from argument 1, except now suppose that each pair of options is presented subsequently. When *A* and *B* are the available options, according to MDoT one should intend to choose *A* over *B*: the egalitarian utilitarian theory gives objective reason to choose *A*, whereas on the person-affecting view there is no objective reason against choosing *A* nor any objective reason against choosing *B* (because in neither case would choosing that option make anyone worse off than they could have been). Similarly, when *B* and *C* are the available options, according to MDoT one should intend to choose *B* over *C*. But, as before, intuitively we should intend to choose *C* over *A*: *C* is much more likely to be right, and there is much more at stake if it is right. So MDoT runs into just the same troubles as DoT did, and for just the same reason.

B. Embrace Cyclicity

Second, either instead of modifying the principle, or in addition to it, one might argue that, sometimes, one ought, rationally, to choose (or intend to choose) in a cyclical manner.¹¹ Against argument 1, one might claim that, in the case I give, one ought to choose *A* over *B*, *B* over *C*, and *C* over *A*. Against argument 2, one might claim that, in the case I give, one ought to choose *A* over *B*, *B* over *C*, and *C* over *A*. One might then develop an account of how one should overcome the difficulties that a rational decision maker with cyclical preferences would face: how such a rational decision maker could avoid money pumps, and so on.

If one were willing to accept the theoretical cost of cyclical preferences or intentions, then one would escape my argument. However, accepting that one's decision theory can lead to cyclical rankings between options is a cost that should be borne only if one has a good reason for accepting such cyclicity. But, despite appearances, Ross's principle is not well motivated.

To see this, we should consider analogues of DoT in other domains. Consider, for example:

Dominance over Times (DoTi).—If, at some times, *x* is better off than *y*, and at no times is *x* worse off than *y*, then *x* is better off than *y* overall.

11. Ross's view, in correspondence, is that one should both modify the principle in the way suggested and accept any resultant cyclicity.

This principle looks *prima facie* compelling. But the principle is false, as can be seen from the following case:

Abraham: (Ω , 1, 2, 3, 4, Ω , Ω , Ω)
Betheneel: (Ω , Ω , 1, 2, 3, 4, 5, Ω)

In this example, the numbers represent how well off the person is at a time, and the omegas represent that the person is not alive at that time. Assume that it is never true to say that someone is better, worse, or equally as well off, at a time, as someone who is not alive at that time.¹² If so, then Abraham is better off at some times than Betheneel is, and is worse off at no times. So DoTi tells us that Abraham is better off overall than Betheneel, even though Betheneel has a greater sum total of well-being, a greater average well-being, a greater peak well-being, and a better end of life than Abraham. So the above example shows that DoTi should be rejected.

As a second example, consider the analogous dominance principle over people.

Dominance over People (DoP).—If, for some people, one population x is better than another population y , and for no one is population x worse than population y , then x is better than y overall.

Again, this principle seems plausible. But only if we assume that one population can't be incomparable with another from the perspective of one person. Consider, for example, the following case:

Population A: (2, Ω , 1)
Population B: (1, 2, Ω)
Population C: (Ω , 1, 2)

In this example, the numbers represent how well off the person is, and the omegas represent that the person is not alive. Assume that it is meaningless to say that someone who exists could be better, worse, or equally as well off in a world in which they do not exist¹³—in which case, according to DoP, population A is better than population B, population B is better than population C, and population C is better than population A. So, according to DoP, A is better than B, B is better than C, and C is better than A. DoP generates a cyclical betterness ranking.¹⁴

12. For a defense of that view, see John Broome, *Weighing Lives* (Oxford: Oxford University Press, 2004), 67.

13. *Ibid.*

14. Note that, though I gave the analogue of argument 1 against DoTi and the analogue of argument 2 against DoP, the analogues of both arguments can be run against both

So we get formally analogous problems if we attempt to use dominance over incomparability in other domains as we do in the case of using dominance over theories. In these other domains, these problems mean that we should reject using dominance over incomparability. We must distinguish carefully between two formulations of dominance principles. For example, if we wish to formulate a dominance principle over people, rather than DoP we should write:

Genuine Dominance over People (GDoP).—If, for some people, one population x is better than another population y , and, for all other people, x and y are equally good, then x is better than y overall.

If we want to formulate a dominance principle over theories, we should write:

Genuine Dominance over Theories (GDoT).—If some theories in which you have credence give you subjective reason to choose x over y , and all other theories in which you have credence give you equal subjective reason to choose x as to choose y , then, rationally, you should choose x over y .

GDoP and DoP appear very similar. But only the former is an acceptable principle of aggregation over people—the latter has apparent plausibility only because of its superficial similarity to the former. In exactly the same way, it is only GDoT that is an acceptable principle of rational choice—DoT's apparent plausibility only arose because it is so easily confused with GDoT. Once this distinction is made, it should be clear that DoT has no theoretical plausibility. So we have no reason to endorse DoT, and therefore we have no reason to endorse the consequent cyclicity that DoT generates.

IV

As a rejoinder, one might motivate MDoT (or DoT) on the grounds of sheer necessity. If we have nonzero credence in nihilism and use expected value reasoning, we get infected, and the expected value of all options is undefined. Using MDoT and accepting the chance of cyclicity is a small price to pay, so one might argue, in order to avoid the expected value of all options being undefined. So it is rational to use a dominance principle like MDoT to reject nihilism.

principles (assuming, at least, that a person can go out of existence, for a time, and come back into existence). To see this, one simply needs to switch round and relabel the examples I give against each principle.

I have sympathy with this line of reasoning. Unfortunately, however, the reasoning does not work. The reason why is because the problem of infectious nihilism is just the tip of a much larger iceberg. The expected value of almost all options is undefined even if we accept MDoT and the resulting cyclicity. Moreover, the reason their expected value is undefined is again because incomparability is infectious under uncertainty. So employing MDoT does not enable us to avoid the infectiousness problem. Moreover, they are undefined for precisely the same reason as before: because of the fact that incomparability is infectious under uncertainty.

Consider, for example, an incomparabilist moral theory according to which there are two values, V_1 and V_2 , but that these values are absolutely incomparable: it is never true, on this theory, that one ought to sacrifice some amount of V_1 for a large enough gain in V_2 , or vice versa. As a result, on this theory, for any two options, x and y , if x involves increasing V_1 and decreasing V_2 , and if y involves increasing V_2 and decreasing V_1 , then x and y are incomparable.¹⁵ The value of x is undefined relative to y .

Now suppose that our decision maker has nonzero credence in this view. If so, then there's trouble, as follows. As noted before, if we take an expectation over possible states of nature, taking the sum, for each state of nature, of the value of that state of nature multiplied by its probability, and the value of one state of nature is undefined, then the expectation as a whole is undefined. So, because, according to this incomparabilist moral view, the value of x relative to y is undefined, for a decision maker with nonzero credence in this moral view, the expected value of x compared with y is undefined, too. This is true even if one has very small credence in the incomparabilist view.

We can go further. For the expected value of x relative to y to be undefined, it need not be the case that option x certainly involves increasing V_1 while decreasing V_2 , and that option y certainly involves increasing V_2 while decreasing V_1 . Just one state of nature in which the value of x is undefined relative to y is sufficient for the expectation of x relative to y to be undefined. So nonzero credence in the incomparabilist moral view combined with the mere empirical possibility that x involves increasing V_1 while decreasing V_2 , and that y involves increasing V_2 while decreasing V_1 , is sufficient for the expected value of x relative to y to be undefined.

Again, this is pretty terrifying. For almost any option and any pair of values, we should have some credence that that option involves increasing one value but decreasing the second value. And the idea that there might exist absolute incomparability between values is not just a specu-

15. There are also other ways in which two options could be incomparable, on this theory. For example, if x involves increasing V_1 and decreasing V_2 , and y involves increasing V_1 by more and decreasing V_2 by less, then x and y are incomparable.

lative hypothesis; it is a definite epistemic possibility.¹⁶ The idea that we should have zero credence in such views is highly implausible—in which case, even if we accept MDoT, if we try to incorporate moral uncertainty into our expected value reasoning, then, on the standard understandings of incomparability and expected value maximization, for every one of us the expected value of almost all options is undefined.

This should make us suspicious of using anything as simple as Rossian dominance-style reasoning in order to reject nihilism. Infectious nihilism is simply one instance of the more general problem of infectious incomparability. In order to understand why we should reject nihilism, we need an understanding of how in general it is rational to act given that one has credence in theories that posit extensive incomparability. If no explanation can be given, then the problem of infectious incomparability gives us grounds for rejecting the entire project of taking moral uncertainty into account in expected value reasoning.

V

In this article I undertook a wholly critical project. I looked at Jacob Ross's argument for the conclusion that we can reject ethical nihilism. I gave two arguments to show that Ross's argument failed. I considered two responses from Ross—that we could modify his principle, and that we could also accept the resultant cyclicity. I argued that these responses also failed. Then, broadening the scope of my argument, I showed that the problem of infectious nihilism is just one instance of the wider problem of infectious incomparability.

Infectious incomparability poses a serious challenge to the project of incorporating moral uncertainty into expected value reasoning. However, I think that it would be too hasty to reject this project, and return to a "my favorite theory" account of decision making under moral uncertainty, or to declare that the idea of there being an "ought" of decision making under moral uncertainty is confused. The problems for such views are too grave. Rather, I think we should try to meet this challenge, and I have some hope that we can. But if we can, nothing as simple as Ross's dominance-style reasoning will suffice. Rather, we need to develop a general account of decision making under uncertainty and incomparability that does not fall prey to the infectiousness problem. As far as I know, no account yet proposed fits this bill.

16. For example, the idea that different people's well-beings are absolutely incomparable is a standard view in classical economics. It is also one way of understanding the "separateness of persons" argument against utilitarianism. Similarly, the idea that very different goods, such as the value of the environment and the value of hedonic states, are absolutely incomparable is not so implausible that one could justifiably assign zero credence to it.